From: AAAI Technical Report FS-92-03. Copyright © 1992, AAAI (www.aaai.org). All rights reserved.

Design of Task-Level Robot Control Systems

Reid Simmons
School of Computer Science / Robotics Institute
Carnegie Mellon University
Pittsburgh, PA 15213
(reids@cs.cmu.edu)

Abstract

Task-level control problems are ones that involve significant coordination of planning, perception and action. Robot systems capable of performing complex tasks are typically composed of many concurrent components. The design problem is to develop a control system that safely and reliably achieves its given tasks. We focus on two aspects of this problem: eliminating unwanted interactions amongs behaviors and dealing with uncertainty arising from incomplete models of the effects of actuators and sensors. Our basic approach to both problems involves combining causal (symbolic) and decision-theoretic reasoning techniques: the causal reasoning help frame the problem, focusing on relevant features, and the decision-theoretic techniques help in making optimal choices.

Introduction

The design of robot control systems has traditionally focused on servo (regulation and tracking) control of dynamic systems characterized by continuous, real-valued variables. For many robot problems, however, *task-level* control is needed. By "task", we mean a clearly defined, operational objective to be achieved using a dynamic control system. Task-level control normally requires the sequencing and coordination of many components, each characterized by discrete, rather than continuous, functions. In task-level control problems, the tasks (goals), environment, and robot hardware (sensors and actuators) are typically orders of magnitude more complex than in traditional servo-control systems.

Task-level control problems are ubiquitous in the real world. Examples include: navigation and control of autonomous mobile robots for applications such as space exploration, disposal of hazardous waste, and remote excavation; rapid fabrication of custom semiconductor integrated circuits and microelectronic assembly; coordination and control of production for flexible parts manufacturing and assembly; on-line tuning of bio-chemical processes; and control of large-scale power systems, especially in cases of emergency.

While sound theoretical design principles exist for servo control [Bollinger and Duffe, 1989], no comparable theory has been enunciated for task-level control. We are interested in providing a firm foundation and computational techniques for automating the design of task-level control systems. The approach is to extend and combine results from AI, decision theory and traditional control theory. We believe that by designing robot systems using such principles, the development cycle can be shortened (due to fewer test/debug cycles) and system safety and reliability can be increased.

An appealing framework for task-level control problems is the *behavior-based* approach [Agre and Chapman, 1987, Brooks, 1986, Brooks, 1989, Connell, 1989]. Intuitively, behaviors are dynamic elements that act over time to collect information and/or affect the environment and robot state. The basic idea in behavior-based control is to construct a small set of generic behaviors that can be parameterized and combined in various ways to perform a wide range of tasks. Different tasks or subtasks can be achieved by switching sets of behaviors at appropriate times. A problem with the behavior-based approach is it assumes that robust primitive behaviors can be developed that act correctly in a wide range of situations. In practice, this can often be very difficult to design *a priori*, due to incomplete, uncertain knowledge about the robot's hardware and environment.

We are investigating a similar methodology, which we call *structured control*, in which deliberative components that handle nominal situations are layered with concurrent reactive behaviors (monitors and error recovery strategies) that provide the necessary safety and reliability [Simmons, 1992a, Simmons, 1992c]. This methodology of integrating deliberative and reactive components provides an engineering basis for developing task-level control systems. The approach accommodates the inevitable incomplete understanding of the task domain by separating the design into the nominal (presumably better understood) behaviors and the infrequently occurring, often less well understood, exceptional situations. By focusing first on the nominal behaviors, it is often easier to design and develop sys-

From: AAAI Technical Report FS-92-03. Copyright © 1992, AAAI (www.aaai.org). All rights reserved.

tems that operate safely and reliably in the majority of situations. This separation of nominal and exceptional behaviors also increases overall system understandability by isolating different concerns: the robot's behavior during normal operation is readily apparent, and strategies for handling exceptions can be developed separately and then layered on to the existing system.

Both behavior-based and structured-control robot systems consist of multiple, concurrent components acting, sensing and planning (to a greater or lesser extent). We are currently investigating two aspects of the design process for such task-level robot control systems: eliminating unwanted interactions between behaviors, and increasing system robustness by explicitly taking uncertainty into account.

Handling Interactions

A primary difficulty in developing both behavior-based and structured-control systems is that behaviors may interact in unexpected ways. To ensure that tasks will be achieved reliably, the designer must detect harmful interactions and prevent them from occurring. While interactions are the bane of most design and planning problems [Simmons, 1992d, Williams, 1991], they are particularly insidious in robot control problems. This is due primarily to the high degree of uncertainty in modeling tasks, environments, and robot hardware and to the real-time aspects of the problem, in particular, the necessity of having components concurrently planning, acting and perceiving (monitoring) the environment.

If the behaviors are too tightly coupled, however, it is very difficult to accurately model all possible interactions. For example, neither the behavior-based nor structured control approaches are really appropriate for aircraft-control maneuvers, due to the complexity of the dynamic interactions [Chiu *et al.*, 1991]. We will limit our investigations, therefore, to domains where the system components, while not independent, are still coupled loosely enough to make the modeling and reasoning problems tractable. Such domains include autonomous mobile robots, process plant control, and many manufacturing operations.

We are extending transformational planning techniques, originally developed for debugging faulty plans [McDermott, 1991, Simmons, 1988, Simmons, 1992d], to deal with interactions in task-level control systems. These techniques use symbolic projection/simulation algorithms to construct causal explanations that describe how, when and why interactions between components (and between the robot and its environment) occur. The planning/debugging techniques analyze the causal explanations to suggest modifications that will eliminate the interactions. Such modifications may involve 1) constraining existing behaviors (e.g., limiting the maximum robot velocity), 2) constraining pair-wise interactions (e.g., sequencing behaviors that were concurrent, or adding resource prioritization constraints), and 3) adding new components (e.g., adding a behavior that removes excess heat caused by mixing chemical reagents in a process plant).

For example, suppose a mobile robot is tasked with delivering objects, keeping the laboratory neat, and maintaining its battery charge. These tasks may arise asynchronously, and the robot system must decide how to prioritize the tasks to avoid resource contention. Rather than using a heuristic or static prioritization scheme, the system can analyze the causal connections between the action preconditions, the system goals (to determine the purposes of the actions) and assumptions underlying the tasks to determine which ordering of tasks is most cost effective. Similar analyses can determine which actions must be taken to transition between tasks effectively (such as setting objects down or fetching particular tools). This is the basis of the work reported in [Goodwin and Simmons, 1992], in which a classical planning algorithm generates alternative plans of action, which are then analyzed according to economic criteria to find the optimal ordering of tasks. In this case, the combination of causal reasoning (planning) and decision-theoretic reasoning (economic analysis) yielded results that neither alone could support.

A major problem in reasoning about interactions is modeling systems with the necessary degree of fidelity, while not making them so complex as to be computationally infeasible. We are investigating the use of the MIDAS simulation package [Simmons, 1991], which integrates multiple representations and reasoning techniques, to model systems and to provide causal explanations for how they behave. MIDAS facilitates the modeling of causal, temporal and spatial aspects of the robot's environment and hardware at various degrees of abstraction (qualitative, semi-quantitative, discrete, and continuous). At this point, we are just beginning to create such models and to determine how the modeling framework must be extended to handle the design of complex robot control systems.

While symbolic, causal reasoning can detect many interactions, it models behaviors at too coarse a level of abstraction to guarantee the reliability of the control system. We intend to augment such reasoning with rigorous quantitative analysis to provide such guarantees. The analysis would be based on algorithms that perform interval-based analyses of control systems modeled as hybrid discrete/continuous networks of behaviors. The basic idea is to propagate the ranges of state variables and event times through the networks and analyze the resulting overlap between various intervals. This approach would essentially extend the control-theoretic algorithms developed for uncertainty propagation and state estimation in continuous dynamic systems with bounded uncertainty (e.g., [Milanese and Belforte, 1982]). It is expected that the semi-quantitative arithmetic reasoning capabilities of MIDAS (which incorporates the Quantity Lattice [Simmons, 1986]) will be well-suited to this reasoning task.

From: AAAI Technical Report FS-92-03. Copyright © 1992, AAAI (www.aaai.org). All rights reserved.

Dealing With Uncertainty

The other aspect of the design problem we are interested in involves reasoning about the uncertainty in one's model of the system. In general, there is significant unpredictability about how a complex robot system will behave, due to incomplete understanding of the effects of actuators and sensors in the real world. Typically, however, the uncertainty in the models can be bounded. In such cases, the control system can often be made more robust by explicitly taking the uncertainty into account [Dean and others, 1990, Kosaka and Kak, 1992, Lozano-Perez *et al.*, 1984].

For example, if one wants a mobile robot to avoid hitting walls, it is more reliable to have it travel in the center of free space, rather than trying to hug the walls. This often conflicts with other goals, such as minimizing travel distance or maintaining accurate knowledge of the robot's current position. In such cases, decision-theoretic techniques are often useful for choosing amongst the various options [Dean *et al.*, 1990].

In a similar vein, we are interested in designing task-specific, selective perception strategies in which the expected cost of sensing is weighed against the expected gain in information [Chrisman, 1992b, Chrisman and Simmons, 1991, Simmons, 1992b]. Sophisticated robot control systems need selective perception for several reasons. First, the number and bandwidth of sensors are limited, but the amount of information potentially available is essentially unlimited. Thus, some method is needed for controlling the information flow. Second, in some situations it is, in fact, more cost-effective to blindly execute a sequence of actions that each have a relatively high probability of success, rather than sensing after every action [Chrisman and Simmons, 1991]. In such situations, selective perception can actually make the system more reliable.

The difficulty is in comparing the various options: determining what is and when it is cost effective to sense. Here again, we expect a combination of causal and decision-theoretic reasoning to be useful. The decision-theoretic reasoning would be used to analyze how sensitive plans are to uncertainty in the environment and action outcomes, how well a particular source of information correlates with the potential success or failure of achieving a goal, etc. Decision-theoretic techniques are computationally quite complex, however, so it is important to focus the reasoning on relevant features of the problem. This is the role for causal, symbolic reasoning: it can *frame* the decision-theoretic problem by pinpointing the important aspects to consider, abstracting away irrelevant details.

The basic idea is to derive causal dependency structures relating the goals (and preconditions) of plans to underlying state variables and parameters of actions. The dependencies then form a framework for focusing the decision-theoretic analysis.

We are investigating a combination of analytical and numeric (Monte Carlo) techniques for characterizing the uncertainty in, and expected cost of, conditional plans. These techniques are used to propagate uncertainty (both symbolic and numerical probability distributions) through the dependency structures (this technique has similarities to the interval-based propagation described in the previous section). Where the variance between the predicted and desired states of the world gets too large, it is an indication that sensing operations may be needed to reduce the variance to acceptable levels. Such decisions about when to sense can often be made at design time. In addition, if the sensing operations are parameterized (such as the direction to point a camera or the sensor resolution to use), this technique can be used to optimize the choice of parameters [Simmons, 1992b].

Suppose, for instance, that a mobile robot has the task of picking up various objects. Due to inaccuracies in its arm and gripper, it must be positioned fairly accurately with respect to an object in order to grasp it successfully. In this case, the definition of "fairly accurately" is very much dependent on the characteristics of the arm and gripper, as well as the relative size of the object, and the uncertainty in the robot's position with respect to that object. Using these analytical techniques, one can determine that for small enough objects, it is optimal for the robot to sense (and correct) its position relative to the object before attempting to pick it up. While the choice of what constitutes "small enough," or at what resolution to sense, really depends on the precise uncertainties involved, the framework of the solution is applicable across robots and objects. Also, while this particular approach relies on an accurate specification of the uncertainty (in terms of probability distribution functions), we believe that it can be adapted to more general models of belief [Chrisman, 1992a].

Conclusions

We have presented several problems relating to the design of task-level robot control systems. These problems stem mainly from interactions between concurrent components and from uncertainty in modeling the system. To address these problems, we advocate using a combination of causal and decision-theoretic reasoning: causal reasoning to focus on relevant aspects of the problem and decision-theoretic reasoning to provide optimal solutions.

Our investigations into designing task-level control problems is just beginning, and it is really too early to know whether useful design tools will emerge. First, it is unclear whether the models of the robot's tasks, environment and hardware can be developed that are tractable, yet provide the necessary information to reason about interactions and uncertainty. Second, it is unclear just how much leverage the causal, symbolic reasoning will provide, and how dependent we will be on more traditional, numeric analysis methods.

We believe, however, that fundamental problems in

From: AAAI Technical Report FS-92-03. Copyright © 1992, AAAI (www.aaai.org). All rights reserved.
reasoning about uncertainty and interactions between behaviors can be tackled using techniques such as those outlined above. The hope is that this will result in principled design techniques for constructing safe and reliable robot control systems.

Acknowledgments

Many thanks to Lonnie Chrisman and Richard Goodwin for their ideas, discussions and implementations. This research is supported by NASA under Contract NAGW-1175.

References

[Agre and Chapman, 1987] Phil Agre and David Chapman. Pengi: An implementation of a theory of activity. In *Proc. National Conference on Artificial Intelligence*, pages 268–272, Seattle, WA, 1987.

[Bollinger and Duffe, 1989] John G. Bollinger and Neil A. Duffe. *Computer Control of Machines and Processes.* Addison-Wesley Publishing Company, 1989.

[Brooks, 1986] Rodney Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2(1), 1986.

[Brooks, 1989] Rodney Brooks. A robot that walks: Emergent behavior from a carefully evolved network. *Neural Computation*, 1(2):253–262, Summer 1989.

[Chiu *et al.*, 1991] S. Chiu, S. Chand, D. Moore, and A. Chaudhary. Fuzzy logic for control of roll and moment for a flexible wing aircraft. *IEEE Control Systems Magazine*, 11(4):42–48, June 1991.

[Chrisman and Simmons, 1991] Lonnie Chrisman and Reid Simmons. Sensible planning: Focusing perceptual attention. In *Proc. National Conference on Artificial Intelligence*, pages 756–761, Los Angeles, CA, July 1991.

[Chrisman, 1992a] Lonnie Chrisman. Abstract probabilistic modeling of action. In *First International Conference on AI Planning Systems*, College Park, MD, June 1992.

[Chrisman, 1992b] Lonnie Chrisman. Planning for closed-loop execution using partially observable markovian decision processes. In *AAAI Spring Symposium on Selective Perception*, Stanford, CA, March 1992.

[Connell, 1989] John Connell. A behavior-based arm controller. *IEEE Journal of Robotics and Automation*, 5(6):784–791, 1989.

[Dean and others, 1990] Tom Dean et al. Coping with uncertainty in a control system for navigation and exploration. In *Proc. National Conference on Artificial Intelligence*, pages 1010–1015, Boston, MA, 1990.

[Dean *et al.*, 1990] Thomas Dean, Kenneth Basye, and Moises Lejter. Planning and active perception. In *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, pages 271–276. DARPA, 1990.

[Goodwin and Simmons, 1992] Richard Goodwin and Reid Simmons. Rational handling of multiple goals for mobile robots. In *First International Conference on AI Planning Systems*, College Park, MD, June 1992.

[Kosaka and Kak, 1992] Akio Kosaka and Avi Kak. Fast vision-guided mobile robot navigation using model-based reasoning and prediction of uncertainties. In *Proc. IEEE International Workshop on Intelligent Robots and Systems*, pages 2177–2186, Raleigh, NC, July 1992.

[Lozano-Perez *et al.*, 1984] Tomas Lozano-Perez, Matthew Mason, and Russell Taylor. Automatic synthesis of fine-motion strategies for robots. *International Journal of Robotics Research*, 3(1), 1984.

[McDermott, 1991] Drew McDermott. Regression planning. *International Journal of Intelligent Systems*, 6:357–416, 1991.

[Milanese and Belforte, 1982] M. Milanese and G. Belforte. Estimation theory and uncertainty intervals evaluation in the presence of unknown but bounded errors: Linear families of models and estimators. *IEEE Trans. on Automatic Control*, 27:408–414, 1982.

[Simmons, 1986] Reid Simmons. Commonsense arithmetic reasoning. In *Proc. AAAI-86*, pages 118–124, Philadelphia, PA, August 1986.

[Simmons, 1988] Reid Simmons. A theory of debugging plans and interpretations. In *Proc. AAAI-88*, pages 94–99, St. Paul, MN, August 1988.

[Simmons, 1991] Reid Simmons. Integrating multiple representations for incremental, causal simulation. In *Proc. Conference on AI, Simulation, and Planning*, pages 88–96, Cocoa Beach, FL, April 1991.

[Simmons, 1992a] Reid Simmons. Concurrent planning and execution for autonomous robots. *IEEE Control Systems*, 12(1):46–50, February 1992.

[Simmons, 1992b] Reid Simmons. Determining sensing resolution and frequency. In *AAAI Spring Symposium on Selective Perception*, Stanford, CA, March 1992.

[Simmons, 1992c] Reid Simmons. Monitoring and error recovery for autonomous walking. In *Proc. IEEE International Workshop on Intelligent Robots and Systems*, pages 1407–1412, July 1992.

[Simmons, 1992d] Reid Simmons. The roles of associational and causal reasoning in problem solving. *Artificial Intelligence*, 53(2–3):159–208, February 1992.

[Williams, 1991] Brian Williams. A theory of interactions: Unifying qualitative and quantitative algebraic reasoning. *Artificial Intelligence*, 51(1–3):39–94, October 1991.