Extended Abstract:
TOWARDS SELF-RELIANT AUTONOMOUS SYSTEMS

Reid Simmons, Sven Koenig, Joaquin Lopez
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Richard Goodwin
IBM Watson Research

To achieve long-term autonomy, autonomous systems must be extremely reliable. We are investigating a number of different techniques to achieve reliability, specifically including the use of probabilistic representations, decision-theoretic planning, and the use of execution monitoring and error recovery strategies that handle situations when plans fail to meet their expectations.

One hallmark of autonomy is the ability to react appropriately to a variety of situations. This presumes that the agent can perceive what situation it is in, and knows the appropriate reactions for that situation. The reason this is difficult for systems that interact with the real world is due to uncertainty, which can arise from uncertainty about the effects of actions, uncertainty in perceiving the environment, and uncertainty in one's model of the environment. One way to handle this is to maintain explicit models of the agent's uncertainty and track when it is outside of the expected bounds of its current plan; Another (orthogonal, but not mutually exclusive) approach is to generate plans that reduce potential uncertainty (e.g., use actions with more certain outcomes, or make sure the agent stays in areas where it has more certain knowledge of the environment).

We have explored both these methods, mainly in the context of mobile robots. In the area of maintaining and tracking explicit models of uncertainty, we make use of Partially Observable Markov Decision Process (POMDP) models to model all three sources of uncertainty [Simmons 95]. POMDP models are automatically compiled from topological maps of the environment, augmented with approximate metric information. As the robot moves, the probability distribution of the POMDP model is updated, using Bayes' rule, to reflect the probable change in robot position. As the robot senses its environment, the observations are used to further update the probability distributions. In this way, while the robot is never exactly sure of its current state, neither does it ever (or rarely) get completely lost. This technique has enabled Xavier, an indoor delivery robot [Simmons 97], to navigate in a natural, peopled environment with greater than 95% success rate in achieving its high-level navigation goals. In more space-relevant applications, we could easily apply these same, or similar, techniques to long-distance navigation of planetary rovers [Krotkov 95] and to tracking the state of autonomous spacecraft (similar to the approach taken by [Williams 96]).

As mentioned, an orthogonal, but compatible, approach is to reduce the potential for uncertainty by judicious choice of plans. This is especially important in situations that are time critical, where replanning or recovering from unexpected situations may not be feasible. Our approach here is to use decision-theoretic planning techniques to create plans that maximize the expected utility of the agent, where utility indicates the agent's preferences (including its risk attitude) [Koenig 96]. The problem here is that the space of such plans is very large. To make planning tractable, the agent needs to focus planning effort on plans that are likely to be of high utility. In particular, we need to determine which plans to focus on, how to refine abstract plans, and when to stop planning and begin execution. This is done by maintaining upper and lower bounds on expected utility, and using a form of sensitivity analysis to rank potential plans and prune out those plans that are guaranteed not to be optimal [Goodwin 95]. This planning approach has been successfully applied to

Xavier, and is also highly applicable to autonomous rovers and spacecraft, especially in dealing with resource-limited situations (such as encounters).

Finally, we have developed an overall architectural framework, based on the Task Control Architecture (TCA), to integrate monitoring and exception handling with task planning and execution. The basic principle, called *structured control*, [Simmons 94a] is that a useful way of developing reliable autonomous systems is to start with behaviors that achieve goals in nominal situations, and then incrementally layer on reactive behaviors (monitors and exception handlers) to handle exceptions. TCA supports the ability to incrementally add monitors and exception handlers to existing systems, making it possible to add robustness to a system without having to change its existing components. This methodology has been applied to several mobile robots [Simmons 92, Simmons 94b], and has demonstrated significant improvements in their abilities to operate reliably for long periods of time.

More recently, we have been investigating how to add additional structure to the monitors and exception handlers to increase their coverage. The idea is to structure the monitors hierarchically, so that high-level monitors are used to detect general categories of exceptional situations (and have associated high-level methods for handling the exceptions), while lower-level monitors are tuned to detect, and handle, more specific situations. For example, in navigating to a given destination, a high-level monitor may checks whether the robot is failing to make progress towards its goal -- if that monitor triggers, the response is to plan an alternate route, or to inform a human if no alternative exists. Lower-level monitors, on the other hand, check for more specific situations (e.g., is the robot trapped in a cul-de-sac? is it getting low on battery voltage?), and likewise trigger more specific recovery strategies (e.g., back up, recharge).

The probabilistic and execution monitoring approaches interact in that the monitors need to know the (likely) state of the world, and need to know what is considered to be "exceptional". The POMDP models can be used for tracking (likely) state; the probabilistic models used by the planner can be used to generate expectations on what should be happening, such as how long the rover should be expected to travel between two points. Integrating this probabilistic information with the monitoring and exception handling approach is still ongoing research. Additional ongoing research addresses the problem of generating such monitors automatically from models of the agent and its environment. This will likely incorporate both probabilistic and symbolic (model-based) reasoning techniques [Williams 96].

In conclusion, autonomous agents cannot blindly execute their plans -- they need extensive capabilities to monitor the state of the world robustly and react appropriately. We are researching several techniques to address this problem, incorporating work in probabilistic reasoning, decision-theoretic planning, architectures for autonomous systems, and model-based reasoning.

References

[Goodwin 95] R. Goodwin. "Sensitivity Analysis for Meta-level Planning", *International Joint Conference on Artificial Intelligence*, Montreal Canada, July 1995.

[Koenig 96] S. Koenig, R. Goodwin and R. Simmons. "Robot Navigation with Markov Models: A Framework for Path Planning and Learning with Limited Computational Resources", Lecture Notes in Artificial Intelligence, *Volume 1093: Reasoning with Uncertainty in Robotics*; L. Dorst, M. van Lambalgen, and R. Voorbraak (Eds.), Springer, pp. 322-337, 1996.

[Krotkov 95] E. Krotkov M. Hebert and R. Simmons. "Stereo Perception and Dead Reckoning for a Prototype Lunar Rover", *Autonomous Robots*, **2:4**, pp. 313-331, 1995.

[Simmons 92] R. Simmons. "Monitoring and Error Recovery for Autonomous Walking" *IEEE International Workshop on Intelligent Robots and Systems*, pp. 1407-1412, July, 1992.

[Simmons 94a] R. Simmons. "Structured Control for Autonomous Robots", *IEEE Transactions on Robotics and Automation*, **10:1**, pp. 34-43, February 1994.

[Simmons 94b] R. Simmons. "Becoming Increasingly Reliable", *Second International Conference on Artificial Intelligence Planning Systems*, Chicago, IL, June 1994.

[Simmons 95] R. Simmons and S. Koenig. "Probabilistic Robot Navigation in Partially Observable Environments", *International Joint Conference on Artificial Intelligence*, Montreal Canada, August 1995.

[Simmons 97] R. Simmons, R. Goodwin, K. Zita Haigh, S. Koenig and J. O'Sullivan. "A Layered Architecture for Office Delivery Robots", *First International Conference on Autonomous Agents*, Marina del Rey, CA, February 1997.

[Williams 96] B. Williams and P. Nayak. "Immobile Robots: AI in the New Millennium", *AI Magazine*, Fall 1996.