Distributed Robust Execution of Qualitative State Plan with Chance Constraints *

Masahiro Ono

Massachusetts Institute of Technology 77 Massachusetts Avenue Cambridge, MA 02139 hiro_ono@mit.edu

Research Objective

Physically grounded AI systems consisting of multiple mobile agents need to conduct complicated tasks cooperatively in dynamic, uncertain environment. Two important capabilities for such systems are robust kinodynamic path planning and distributed plan execution on a hybrid discrete/continuous plant. For example, a fleet of autonomous underwater vehicles (AUV) shown in Figure 1, which conducts scientific observations cooperatively for up to 20 hours without human supervision, should ideally navigate themselves to areas of scientific interest according to a game plan provided by scientists.

Our plan formalism is Qualitative State Plan (QSP) (Léauté 2005), which specifies the desired evolution of the qualitative state of the system as well as the flexible temporal constraints. This approach elevates the interaction between the human operator and the robotic system, to a more abstract level where the operator is able to qualitatively command the tasks. A centralized model-based QSP executive called Sulu (Léauté 2005) generates optimal path and schedule that is consistent with a given QSP in deterministic environment.

Real-world systems, however, are exposed to stochastic disturbances. Stochastic systems typically have a risk of failure due to unexpected events, such as unpredictable tides and currents that affect the AUV's motion. AUV operators want to limit the risk of losing AUV by colliding with seafloor. Thus the kinodynamic path planning has to be robust in existence of disturbance.

The plan executive should ideally be distributed for a number of reasons. First, in many cases such as underwater, the inter-vehicle communication is limited. Second, the leader vehicle that has the centralized plan executive is the single point failure. Third, computation burden concentrates in the leader vehicle. Distributed plan executive makes system more robust and efficient.

My research objective is to develop a *distributed* modelbased QSP executive that is *robust* in a stochastic environment.



Figure 1: Autonomous underwater vehicle *Dorado* of Monterey Bay Aquarium Research Institute, conducting bathymetric mapping mission

Research Plan

There are four major steps to develop distributed robust QSP executive.

- Step 1: Develop robust kinodynamic path planner
- Step 2: Develop centralized robust QSP executive
- Step 3: Develop distributed robust QSP executive
- Step 4: Experiment on real-world system.

Step 1 has been achieved, which will be introduced in the next section. The robust path planner developed in Step 1 will be merged to Sulu, a centralized model-based QSP executive, to achieve Step 2. To proceed from Step 2 to Step 3, two problems must be solved; distribution of QSP and collision avoidance. For Step 4, we are collaborating with Monterey Bay Aquarium Research Institute to test the algorithms using their AUV.

Robust Kinodynamic Path Planner

The task of robust kinodynamic path planner is to find the control sequence (action sequence) of a stochastic system that maximizes performance while guaranteeing that the probability of failure is less than the given upper bound (chance constraint). We developed a new efficient robust kinodynamic path planner called Bi-stage Robust Motion Planning (BRMP) algorithm (Ono & Williams 2008). The key notion is *risk allocation*, which is explained in the next subsection.

Racing Car Example Imagine a racing car example shown in Figure 2. The dynamics of the vehicle is stochastic and the distribution of uncertainty is unbounded. The task is to plan a control sequence of wheel and acceleration

^{*}This research is funded by The Boeing Company grant MIT-BA-GTA-1

Copyright © 2008, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.



Figure 2: Risk allocation strategies on the racing car example

that minimizes the time to reach a goal, with the guarantee that the probability of crashing into a wall during the race is less than a certain probability, say, 0.1% (*chance constraint*). Planning the control sequence is equivalent to planning the nominal path, which is shown as the solid lines in the Figure 2. To limit the probability of crashing into the wall, a good driver would set the safety margin, which is colored in dark gray in Figure 2, and then plan the nominal path out of the safety margin. In other words, the driver *tightened* the original constraints (the walls) and set new constraints on the nominal path, which is shown as the dotted line.

The driver wants to set the safety margin as small as possible to make the nominal path shorter. However, since the probability of crash during the race is bounded, there is a certain lower bound on the size of the safety margin. We assume here that the total area of the safety margin is lowerbounded. Given this constraint, there are different strategies of setting a safety margin; in Figure 2(a) the width of the margin is uniform; in Figure 2(b) the safety margin is narrow around the corner, and wide at the other places.

An intelligent driver would take the strategy of (b), since he knows that going closer to the wall at the corner is effective to make the path shorter while doing so at the straight line is not. A key observation here is that taking a risk (i.e. setting narrow safety margin) at the corner results in a greater reward (i.e. time saving) than taking the same risk at the straight line. This gives rise to the notion of *risk allocation*. The good risk allocation strategy is to save risk when the reward is small while taking it when the reward is great.

Another important observation is, once risk is allocated and the safety margin is fixed, the stochastic path planning problem with chance constraint has been reduced to a deterministic nominal path planning problem with tightened constraints. This can be solved quickly with existing deterministic path planning algorithms.

These two observations lead to bi-stage optimization algorithm as shown in Figure 3, in which its upper stage allocates risk to each time step while its lower stage tightens constraints according to the risk allocation and solves the resulting deterministic problem. We call this algorithm Bistage Robust Motion Planning (BRMP) (Ono & Williams 2008).



Figure 3: Architecture of Bi-stage Robust Motion Planning

Table 1: Performance comparison on the AUV depth navigation problem with chance constraint $P_{Fail} \le 0.05$. (a) ER: Ellipsoid relaxation approach, (b) BRMP: Bi-stage Robust Motion Planning, PC: Particle Control

Algorithm used	(a) ER	(b) BRMP	(c) PC
Resulting P _{Fail}	$< 10^{-5}$	0.037	0.085
Average altitude [m]	99.3	55.2	51.1
Computation time [sec]	1.9	4.1	481.2

Implementation BRMP is implemented and tested on a AUV depth navigation case. The task is to minimize AUV's altitude from sea floor while limiting the probability of crashing into it. The dynamics model is taken from the actual AUV developed by Monterey Bay Aquarium Research Institute (Figure 1), and the actual bathymetric data of the Monterey Bay is used. The deterministic planning algorithm used in the lower-stage has been demonstrated in the actual AUV mission.

The simulation result is shown in Table 1, along with the two prior arts; ellipsoidal relaxation approach (van Hessem 2004) and Particle Control (Blackmore 2006). It is shown from this result that BRMP is much less suboptimal than ellipsoidal relaxation approach while achieving substantial speed up from Particle Control.

References

Blackmore, L. 2006. A probabilistic particle control approach to optimal, robust predictive control. In *Proceedings of the AIAA Guidance, Navigation and Control Conference*.

Léauté, T. 2005. Coordinating agile systems through the model-based execution of temporal plans. Master's thesis, Massachusetts Institute of Technology.

Ono, M., and Williams, B. C. 2008. An efficient motion planning algorithm for stochastic dynamic systems with constraints on probability of failure. In *Proceedings* of the Twentieth National Conference on Artificial Intelligence (AAAI).

van Hessem, D. H. 2004. *Stochastic inequality constrained closed-loop model predictive control with application to chemical process operation*. Ph.D. Dissertation, Delft University of Technology.