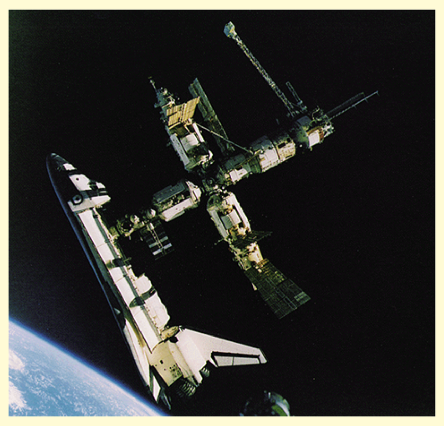

Computing Probabilistic Guarantees of Safety

Kendra Lesser, Meeko Oishi
lesser@unm.edu, oishi@unm.edu

Problem

Stochastic reachability provides a theoretical framework for generating probabilistic guarantees of safety in cyber-physical systems. Unfortunately, its applicability to realistic (i.e. high-dimensional and partially observable) systems is currently limited, because of a lack of feasible computation strategies.

Research Challenges

We need to develop computational methods that are:

- Able to calculate safety or *viability* probabilities.
- Able to synthesize safety-maximizing controllers.
- Extendable to higher-dimensional systems.
- Applicable to systems with noisy or incomplete state observations.

Contributions

1. A numerical method to approximate reachable sets and optimal controllers for partially observable stochastic hybrid systems, based on point-based value iteration.

2. A particle approximation to estimate reachable sets using mixed-integer linear programming, which is applicable to systems with more than three dimensions.

References

[1] A. Abate, M. Prandini, J. Lygeros, and S. Sastry, "Probabilistic reachability and safety for controlled discrete time stochastic hybrid systems," *Automatica*, vol. 44, no. 11, pp. 2724–2734, 2008.

[2] K. Lesser, M. Oishi. and R. S. Erwin, "Stochastic reachability for control of spacecraft relative motion," in Proceedings of the *IEEE Conference on Decision and Control*, 2013.

[3] K. Lesser and M. Oishi, "Computational Techniques for Reachability Analysis of Partially Observable Discrete Time Stochastic Hybrid Systems," *IEEE Transactions on Automatic Control*, 2014. Submitted, under review.

[4] K. Lesser and M. Oishi, "Reachability for Partially Observable Discrete Time Stochastic Hybrid Systems," 2014, http://arxiv.org/abs/1401.5756.

[5] J. M. Porta, N. Vlassis, M. T. Spain, and P. Poupart, "Point-based value iteration for continuous POMDPs," *Journal of Machine Learning Research*, vol. 7, pp. 2329–2367, 2006.

Partially Observable Systems

In the perfectly observed case, stochastic reachability can be formulated as a multiplicative cost optimal control problem [1]:

$$\max \mathbb{E}\left[\prod_{i=1}^{N} \mathbf{1}_K(q_i, x_i) \middle| q_0, x_0\right]$$

Approach: Assuming that the controller does not have access to the true state, we would like to apply results from additive cost partially observable Markov decision processes (POMDPs).

1. Reduce the problem to one with perfect information by using a sufficient statistic.

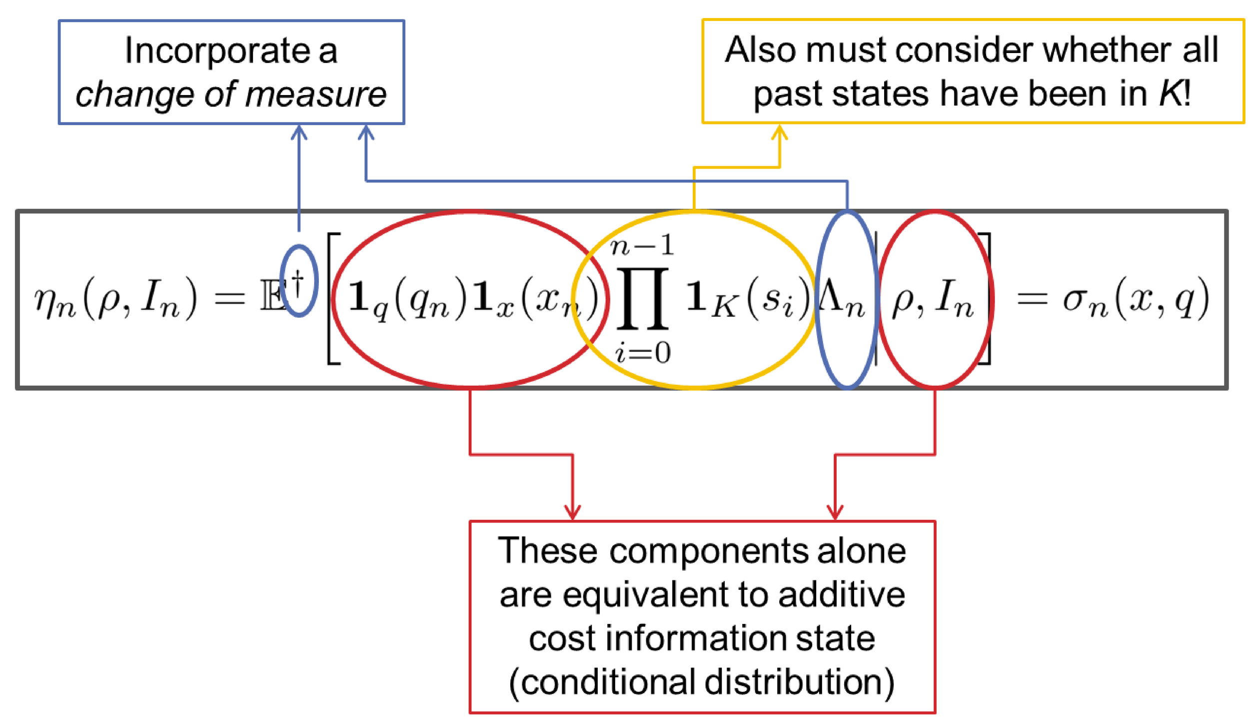

2. The value function can be represented by a finite set of α-functions, because it is piecewise-linear and convex [3] (as in the case of additive cost POMDPs).

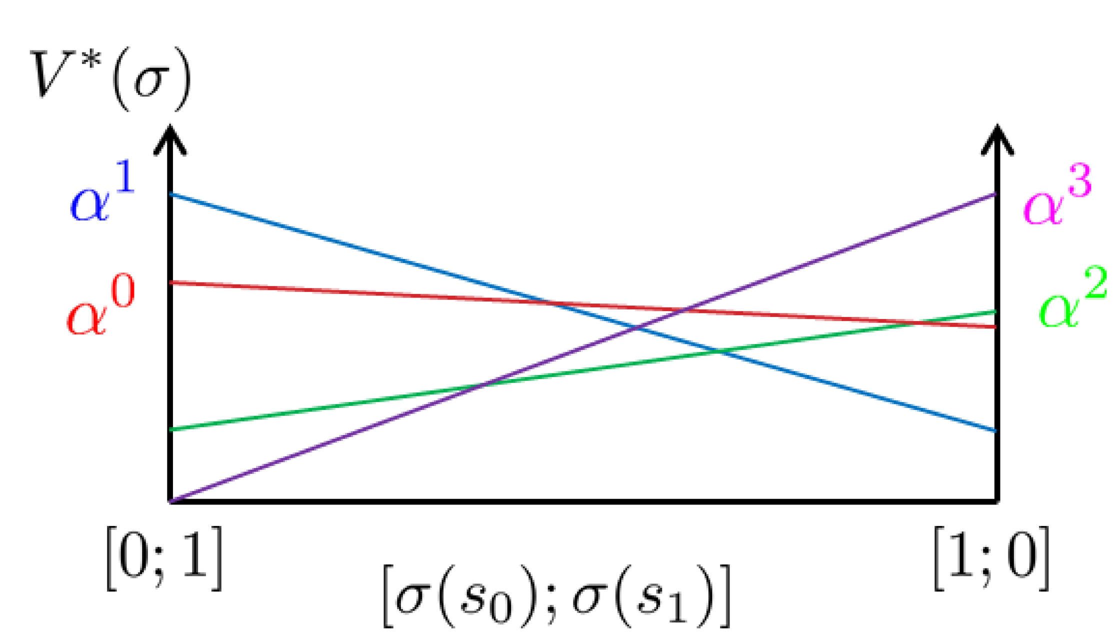

Stochastic Reachability

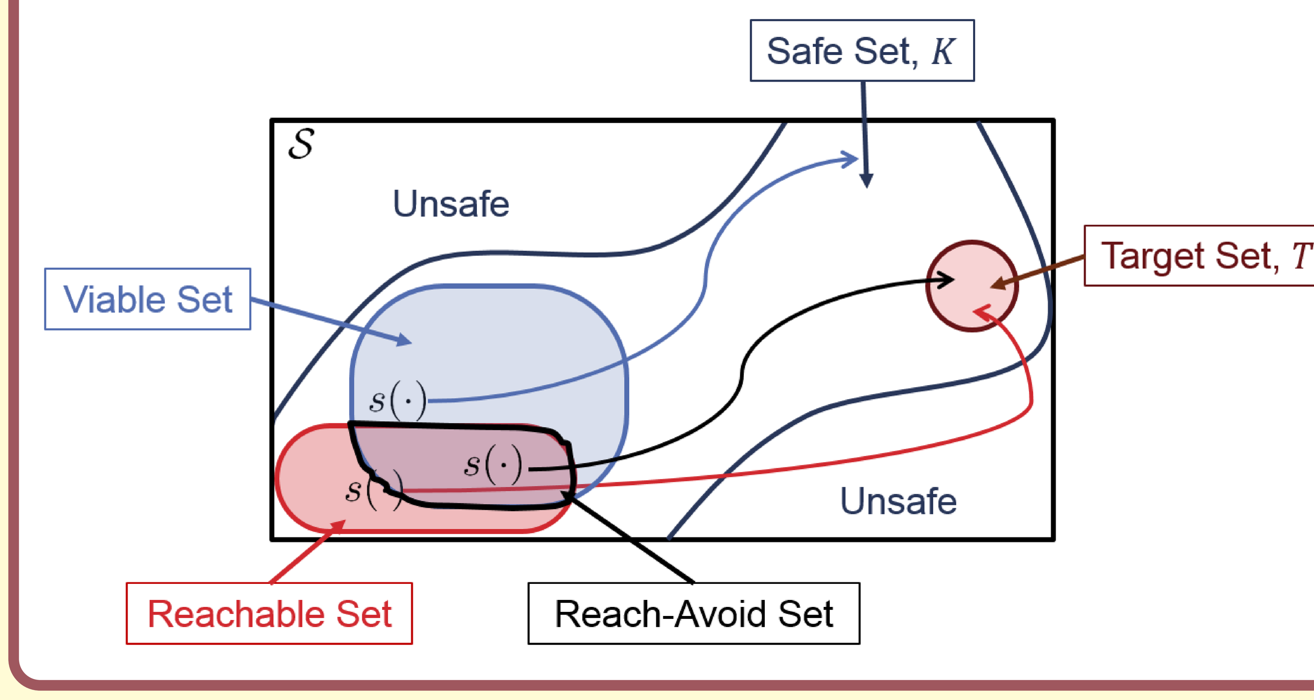

Application: Temp. Regulation

- Controller only has available a noisy measurement of the temperature, but wants to keep a room within $K = [17.5, 22]$ degrees Celsius.
- As a proof of concept, we apply a version of point-based value iteration similar to [5].

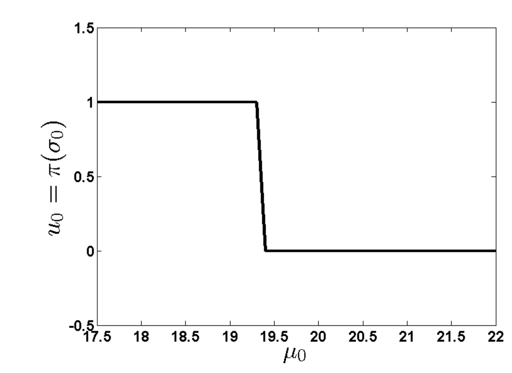

- α-functions produce value function estimate and encode optimal control actions.

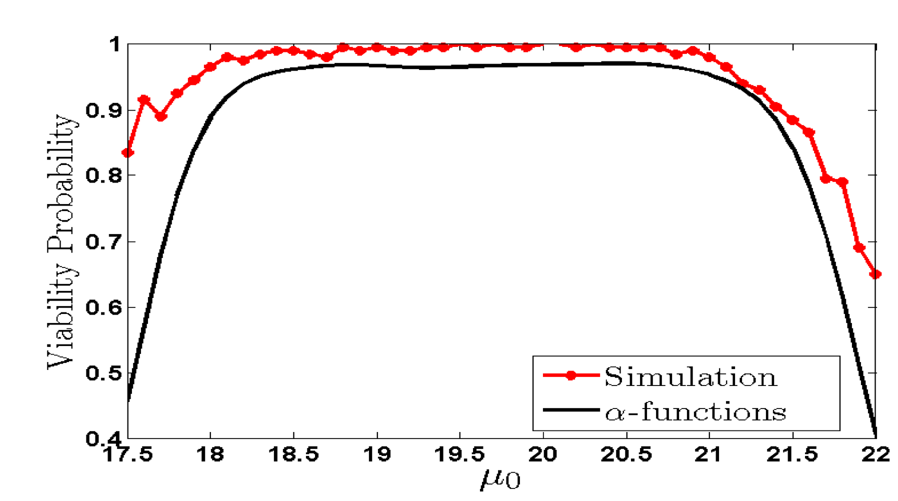

- Comparing α-function approximation to simulation with α-function policy, α-functions lower bound true value function, as expected.
- Use Gaussian sums to represent the α-functions and information states σ, which requires a Gaussian sum approximation to the indicator function.

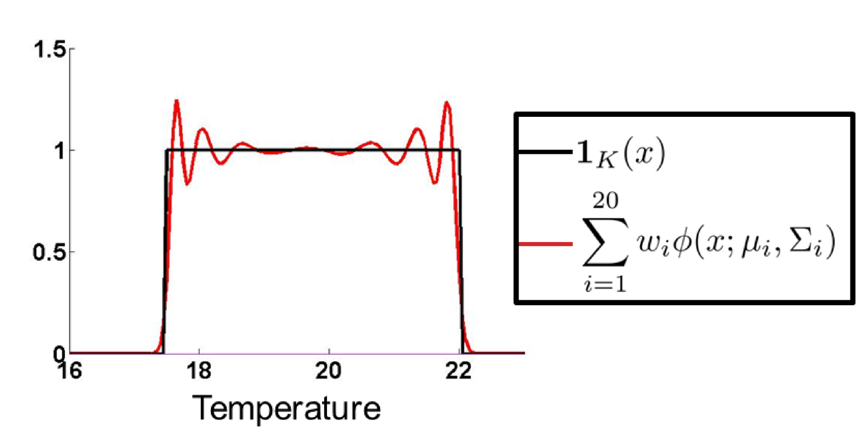

Stochastic Reachability for Spacecraft Rendezvous

With four-dimensional in-plane linearized relative dynamics for rendezvous between two spacecraft, we can sample from the stochastic noise to create particle trajectories for the dynamics. We then use mixed-integer linear programming to find control inputs that maximize the number of sample trajectories satisfying reach-avoid constraints [2].

$$\max_{U\in\mathcal{U}} \frac{1}{M}\sum_{i=1}^{M}\left[\left(\prod_{j=1}^{N-1}\mathbf{1}_K(X_j^{(i)})\right)\mathbf{1}_T(X_N^{(i)})\right]$$

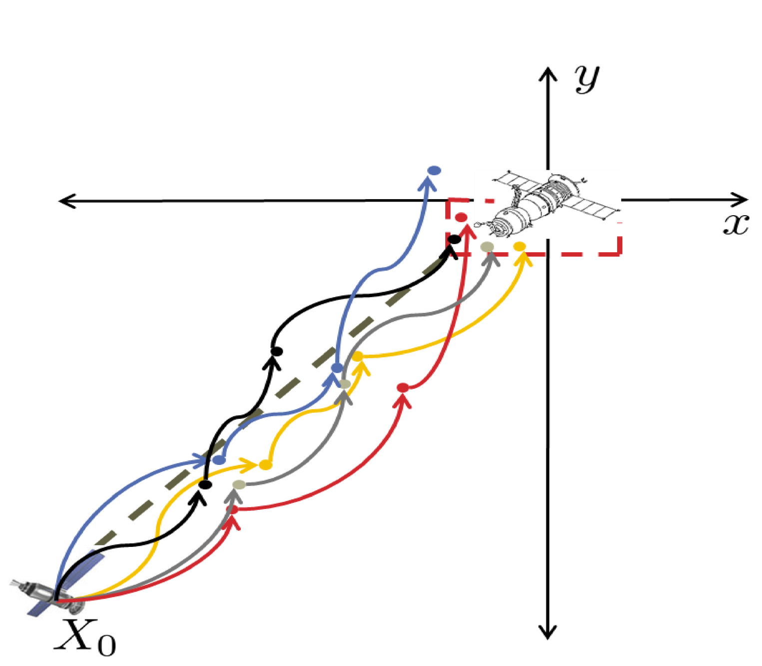

Sample trajectories

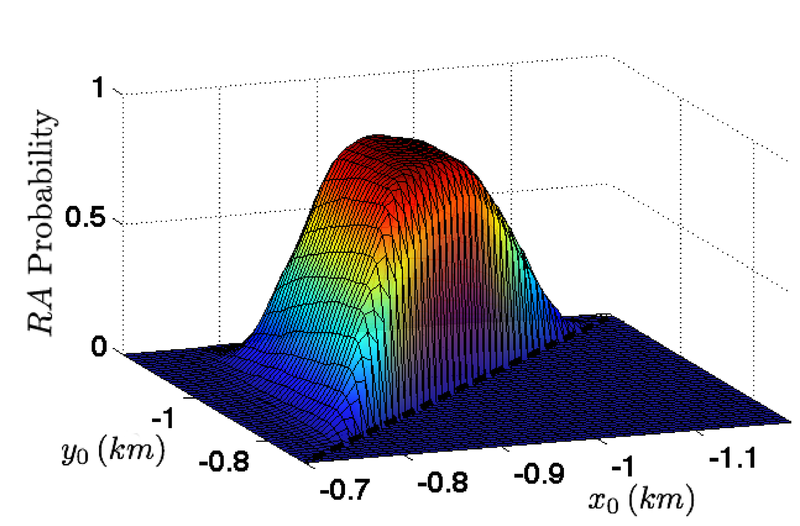

Reach-avoid probabilities for varying initial position and fixed initial velocity

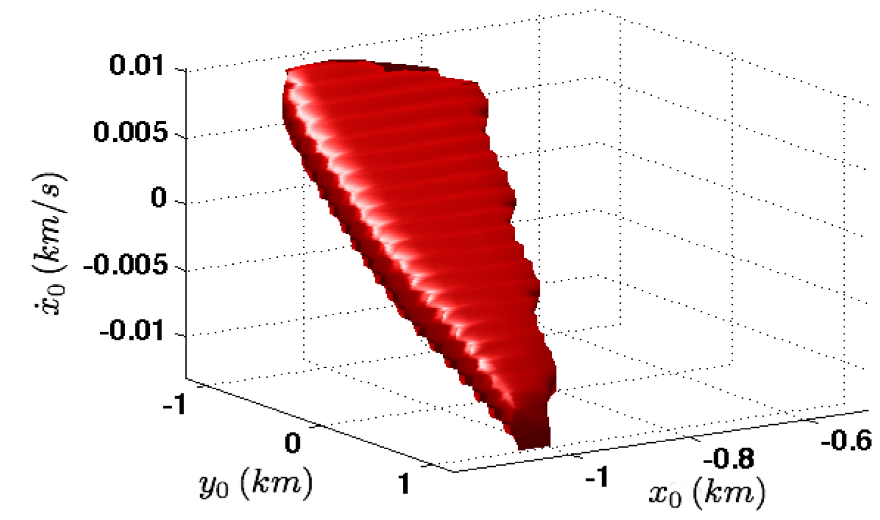

Reach-avoid probabilities for varying initial position and x-velocity, fixed y-velocity

Acknowledgements

This research was funded by National Science Foundation (NSF) Career Award CMMI-1254990, NSF Award CPS-1329878, NSA Science of Security Lablet at North Carolina State University (subaward to the University of New Mexico), the Air Force Research Laboratories (AFRL) Space Scholars Program, and an AFRL Summer Faculty Fellowship.