Asymptotic Optimality in Sampling-based Motion Planning

Sertac Karaman Emilio Frazzoli

Although one of the fundamental problems in robotics, the motion planning problem is inherently hard from a computational point of view. In particular, the piano movers' problem [1], [2] is known to be PSPACE-hard, which implies that any algorithm aimed to solve this problem (with completeness guarantees) is expected not to scale well with increasing number of dimensions of the configuration space. In fact, the computation time required by well-known complete algorithms scale exponentially with the dimensionality of the configuration space in the worst case [3], which makes them impractical, e.g., in motion planning problems for robotic arms with several joints.

The optimal motion planning problem is known to be significantly harder computationally when compared to finding just a feasible solution, even when the dimensionality of the configuration space is fixed. In particular, it is known that the optimal motion planning problem is NP-hard for a point mass moving among polygonal obstacles in a three dimensional configuration space [4]. Early research on motion planning during 1980s has involved optimal motion planning algorithms only with limited success [5], [6], [7]. However, especially after the emergence of a deeper understanding of the computational complexity of the problem during the late 1980s, the community shifted towards designing algorithms that can quickly find feasible solutions, usually with no optimality guarantees.

Among others, sampling-based algorithms, first proposed by Kavraki et al. [8] in the mid-1990s under the name of Probabilistic RoadMaps (PRMs), have become one of the first practical approaches to address problems that involve high-dimensional configuration spaces. A few years later, LaValle and Kuffner [9] have focused on online settings and systems with differential constraints, and proposed the Rapidly-exploring Random Tree (RRT) algorithm, which is arguably one of the most widely-used motion planning algorithms as of today. Both the PRM and the RRT algorithms are shown to be *probabilistically complete* in the sense that, loosely speaking, the probability that the algorithm returns a solution, if one exists, converges to one as the number of samples approaches infinity; In fact, the probability that these algorithms fail to return to a feasible solution, when one exists, converges to zero exponentially fast [9], [10]. Beyond, their theoretical guarantees, these planners have also been shown to perform well in practice. In particular, planners based on the RRT have been implemented on various robotic platforms and showcased in major robotics events [11].

The authors are with the Laboratory of Information and Decision Systems, Massachusetts Institute of Technology.

During the last decade, sampling-based algorithms have benefited from the increasing computational power embedded in commodity personal computers, allowing these algorithms to find feasible solutions to challenging problem instances within very short execution times. Encouraged by the success of sampling-based planners, recently many researchers have started revisiting the optimal motion planning problem, this time in the context of sampling-based algorithms [12], [13], [14].

Non-optimality of the trajectories returned by the RRT, however, was long known to practitioners. Indeed, LaValle and Kuffner [9] note in their seminal paper describing the RRT that "it is obvious that the generated trajectories are not optimal, even within their path (homotopy) class." This empirical observation has been formalized very recently by Karaman and Frazzoli [14] who have shown that the probability that the RRT converges to an optimal solution is in fact zero. In the same paper, Karaman and Frazzoli have also introduced the concept of *asymptotic optimality*, which amounts to almost-sure convergence to optimal solutions with increasing number of samples. They have also proposed a new algorithm, called the RRT*, with the asymptotic optimality property. Surprisingly, the asymptotic computational complexity of the RRT* algorithm is the same as that of the RRT algorithm. That is, the RRT* algorithm achieves asymptotic optimality, which the RRT algorithm lacked, with no substantial computational overhead when compared to the RRT. See Figure 1 for a comparison of the two algorithms in an illustrative example. The authors have also introduced an efficient asymptotically-optimal variant of PRM, called the PRM*, and an incremental version of the PRM*, called the Rapidly-exploring Random Graph (RRG) algorithm. In their paper, Karaman and Frazzoli [14] also establish concrete links between sampling-based motion planning algorithms and the theory of random geometric graphs.

From a practical perspective, a salient feature of asymptotically-optimal incremental sampling-based algorithms, such as the RRT*, is their anytime flavor: the algorithm provides a feasible solution quickly and improves this solution towards an optimal one if allowed more computation time. The RRT* algorithm has recently been implemented on various robotic platforms including robotic cars [15] and manipulation platforms [16] that feature up to fourteen-dimensional configurations spaces. The results presented in these references are illustrated in Figures 2 and 3. Experimental results indicate that the difference in computation time between the RRT and the RRT* is not very large, e.g., less than a factor of two, although the quality of the final solution differs significantly [16]. Moreover, it was shown by Bialkowski et al. [17] that this difference in computation time can be marginalized by intelligent use of massive parallelization on dedicated parallel computing hardware such as general purpose Graphics Processing Units (GPUs). More recently, the RRT* algorithm was implemented for complex dynamical systems such as race cars, and it was shown that the trajectories generated by the RRT* algorithm resemble those used by expert race car drivers [18].

Most importantly, the algorithms proposed by Karaman and Frazzoli [14] have made it possible to attack novel classes of problems that could not be handled using incremental sampling-based algorithms, such as the RRT, before. For instance, the RRT* algorithm was used to solve a class of pursuit-evasion games with probabilistic completeness and probabilistic soundness guarantees [19], while the RRG algorithm was used to solve motion planning problems with complex task specifications given in the form of deterministic μ-calculus [20]. These new algorithms also inherit the anytime flavor, which allows practical real-time implementations.

Looking forward, several indicators are present to predict that the next decade will experience another significant increase in the amount of affordable computational power, which will likely be based on massively parallel computational architectures such as general purpose GPUs. Being amenable to parallel implementations [21], [17], sampling-based algorithms are likely to largely benefit from this trend. Hence, it is likely that a significant research effort will be devoted to asymptotically-optimal motion planning using sampling-based algorithms as well as to addressing novel problems in an efficient manner using such algorithms.

This poster presentation outlines fundamental results, open problems, and conjectures regarding asymptotic optimality of sampling-based algorithms, and presents novel problem domains that can be handled by variants of the new algorithms. Fundamental results include a through analysis of several sampling-based algorithms from the perspective of asymptotic optimality and computational complexity. Open problems include convergence rates and lower bounds, which we present partial results on. New problem domains include differential games, planing with complex task specifications, and also novel planning under uncertainty problems that involve continuous-time continuous-space models with Brownian process noise. Finally, the presentation also points out novel connections between sampling-based motion planning and the theory of random geometric graphs.

REFERENCES

[1] J. T. Schwartz and M. Sharir. On the 'piano movers' problem: II. general techniques for computing topological properties of real algebraic manifolds. *Advances in Applied Mathematics*, 4:298–351, 1983.

[2] J.E. Hopcroft, J.T. Schwartz, and M. Sharir, editors. *Planning, Geometry, and Computational Complexity of Robot Motion*. Ablex, Norwood, 1987.

[3] J. Canny. *The Complexity of Robot Motion Planning*. MIT Press, 1987.

[4] J. Canny and J.H. Reif. New lower bound techniques for robot motion planning problems. In *Proceedings of the 28th Annual IEEE Symposium on Foundations of Computer Science*, 1987.

[5] C. Papadimitriou. An algorithm for shortest-path motion in three dimensions. *Information Processing Letters*, 20:259–263, 1985.

[6] M. Sharir and A. Schorr. On shortest paths in polyhedral spaces. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, 1984.

[7] J. Reif and J.A. Storer. Shortest paths in euclidean space with polyhedral obstacles. Technical report, Harvard University, 1985.

[8] L. E. Kavraki, P. Svestka, J. C. Latombe, and M. H. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Transactions on Robotics and Automation*, 12(4):566–580, 1996.

[9] S. LaValle and J. Kuffner. Randomized kinodynamic planning. *International Journal of Robotics Research*, 20(5):378–400, May 2001.

[10] L. E. Kavraki, M. N. Kolountzakis, and J. C. Latombe. Analysis of probabilistic roadmaps for path planning. *IEEE Transactions on Roborics and Automation*, 14(1):166–171, 1998.

[11] Y. Kuwata, J. Teo, G. Fiore, S. Karaman, E. Frazzoli, and J.P. How. Real-time motion planning with applications to autonomous urban driving. *IEEE Transactions on Control Systems*, 17(5):1105–1118, 2009.

[12] B. Raveh, A. Enosh, and D. Halperin. A little more, a lot better: Improving the path quality by a path-merging algorithm. *IEEE Transactions on Robotics*, 27(2):365–371, April 2011.

[13] L. Jaillet, J. Cortes, and T. Simeon. Sampling-based path planning on configuration-space costmaps. *IEEE Transactions on Robotics*, 26(4):635–646, August 2010.

[14] S. Karaman and E. Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7):846–894, June 2011.

[15] S. Karaman, M. Walter, A. Perez, E. Frazzoli, and S. Teller. Anytime motion planning using the RRT*. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2011.

[16] A. Perez, S. Karaman, M. Walter, A. Shkolnik, E. Frazzoli, and S. Teller. Asymptotically-optimal path planning for manipulation using incremental sampling-based algorithms. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (to appear)*, 2011.

[17] J. Bialkowski, S. Karaman, and E. Frazzoli. Massively parallelizing the RRT and the RRT*. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (to appear)*, 2011.

[18] J. Jeon, S. Karaman, and E. Frazzoli. Anytime computation of time-optimal off-road vehicle maneuvers using the RRT*. In *IEEE Conference on Decision and Control (CDC) (to appear)*, 2011.

[19] S. Karaman and E. Frazzoli. Incremental sampling-based algorithms for a class of pursuit-evasion games. In *Workshop on Algorithmic Foundations of Robotics (WAFR)*, Singapore, December 2010.

[20] S. Karaman and E. Frazzoli. Sampling-based motion planning with deterministic μ-calculus specifications. In *IEEE Conference on Decision and Control (CDC)*, Shanghai, China, December 2009.

[21] N. Amato and L. Dale. Probabilistic roadmap methods are embarrassingly parallel. In *IEEE Conference on Robotics and Automation (ICRA)*, 1999.

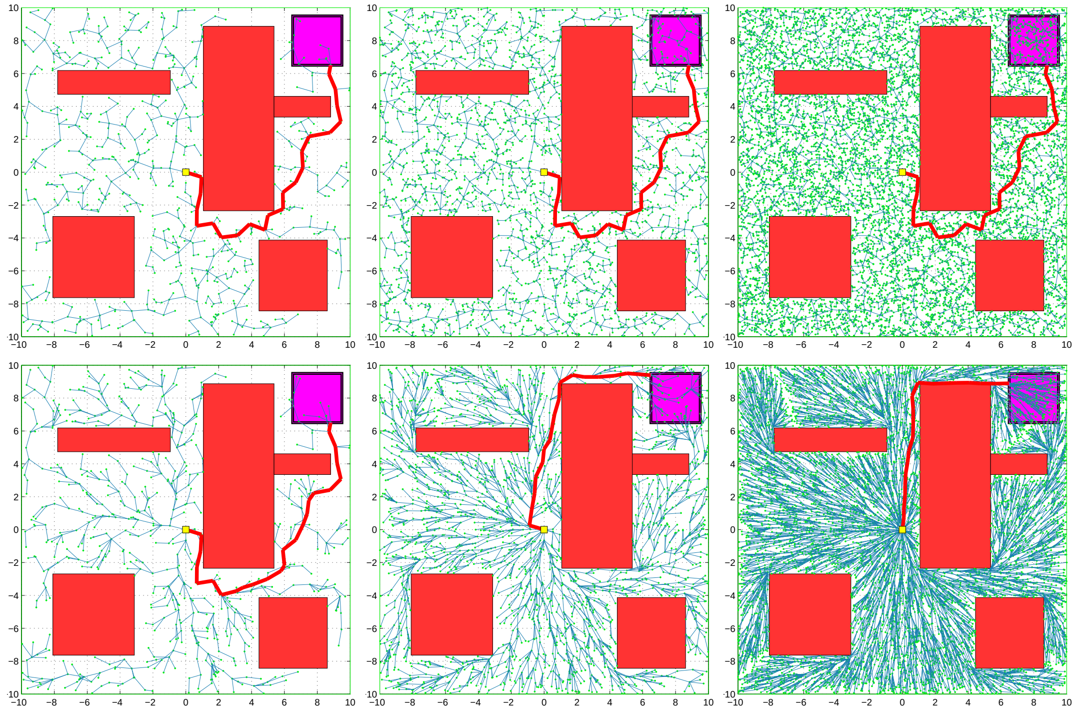

Fig. 1. A comparison of the RRT and RRT* algorithms. The obstacles are shown in red, the tree is shown in blue, the goal region is shown in magenta, and the minimum-cost path is highlighted with red. The top three figures show the RRT algorithm at iterations 1000, 3000, and 10000. The bottom three figures show the RRT* algorithm at the same stages.

(a) (b) (c)

Fig. 2. Paths traversed by the RRT and the RRT* algorithms in a real-time planning scenario for an autonomous car-like robot are shown in Figures (a) and (b), respectively. The obstacles are shown in red, and the goal region is shown in green. The robot is shown in Figure (c).

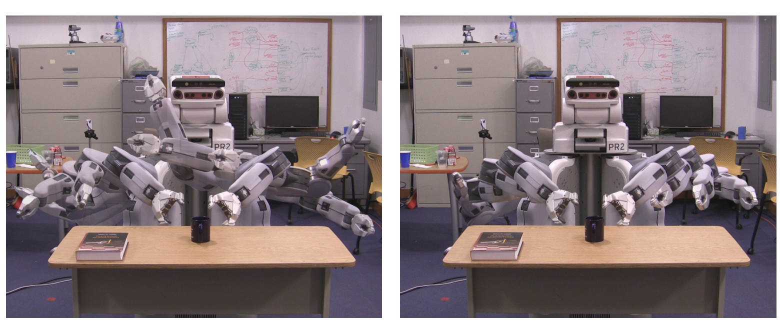

(a) (b)

Fig. 3. Path planning for the PR2 robot for both arms (12-dimensional configuration space). The plans generated by the RRT and the RRT* algorithm are illustrated in Figures (a) and (b), respectively. The robot starts with both arms under the table; the goal is to reach the pre-grasp pose in which the hands point towards the cup.