Informed Steiner Trees: Sampling and Pruning for Multi-Goal Path Finding in High Dimensions (Extended Abstract)*

Nikhil Chandak,[1 **] **Kenny Chour,**[2 **] **Sivakumar Rathinam,**[2] **R. Ravi**[3]

[1] IIIT-Hyderabad
[2] Texas A&M University, College Station
[3] Carnegie Mellon University, Pittsburgh
nikhil.chandak@research.iiit.ac.in, {ckennyc, srathinam}@tamu.edu, ravi@andrew.cmu.edu

Abstract

We interleave sampling based motion planning methods with pruning ideas from minimum spanning tree algorithms to develop a new approach for solving a Multi-Goal Path Finding (MGPF) problem in high dimensional spaces. The approach alternates between sampling points from selected regions in the search space and de-emphasizing regions that may not lead to good solutions for MGPF. Our approach provides an asymptotic, 2-approximation guarantee for MGPF. We also present extensive numerical results to illustrate the advantages of our proposed approach over uniform sampling in terms of the quality of the solutions found and computation speed.

Introduction

Multi-Goal Path Finding (MGPF) problems aim to find a least-cost path for a robot to travel from an origin (s) to a destination (d) such that the path visits each node in a given set of goals ($\bar{T}$) at least once. In the process of finding a least-cost path, MGPF algorithms also find an optimal sequence in which the goals must be visited. When the search space is discrete ($i.e.$, a finite graph), the cost of traveling between any two nodes can be computed using an all-pairs shortest paths algorithm. In this case, the MGPF encodes a variant of the Steiner[1] Traveling Salesman Problem (TSP) and is NP-Hard (Kou, Markowsky, and Berman 1981). In the general case, the search space is continuous and the least cost to travel between any two nodes is not known a-priori. This least-cost path computation between any two nodes in the presence of obstacles, in itself, is one of the most widely studied problems in robot motion planning (Kavraki et al. 1996; Kuffner and LaValle 2000). We address the general case of MGPF as it naturally arises in active perception (Best, Faigl, and Fitch 2016; McMahon and Plaku 2015), surface inspection (Edelkamp, Secim, and Plaku 2017) and logistical applications (Janoš, Vonásek, and Pěnička 2021; Otto et al. 2018; Macharet and Campos 2018). MGPF is notoriously hard as it combines the challenges in Steiner TSP and the least-cost path computations in the presence of obstacles; hence, we are interested in finding approximate solutions for MGPF.

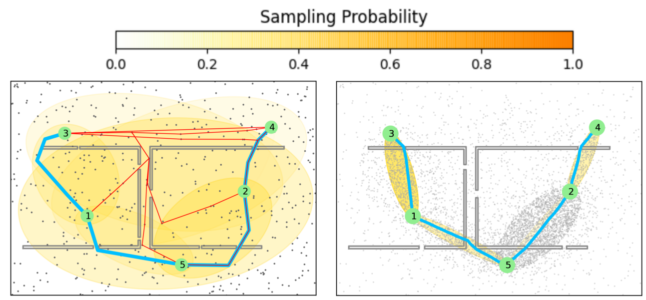

Figure 1: The Steiner tree (blue lines) computed from IST*, showing the advantage of pruning. The environment has 2 U-shaped obstacles with open channels. Numbered nodes are terminals to be spanned. Ellipses are an edge's informed set, with color intensity related to sampling probability. Sampled points in free space are part of the roadmap (shown in grey). **Left**: without pruning, all non-MST edges (thin-red lines) are also sampled. **Right**: after pruning, the roadmap is only densified around edges that can lead to an optimal MST faster.

Irrespective of whether the search space is discrete or continuous, Steiner trees spanning the origin, goals and the destination play a critical role in the development of approximation algorithms for MGPF. In the discrete case, doubling the edges in a suitable Steiner tree, and finding a feasible path in the resulting Eulerian graph leads to 2-approximation algorithms for MGPF (Kou, Markowsky, and Berman 1981; Mehlhorn 1988; Chour, Rathinam, and Ravi 2021). This approach doesn't readily extend to the continuous case because we do not a-priori know the travel cost between any two nodes in $T := \{s, t\} \bigcup \bar{T}$. One can appeal to the well-known sampling-based methods (Karaman and Frazzoli 2011; Kavraki et al. 1996; Gammell, Srinivasa, and Barfoot 2014, 2015) to estimate the costs between the nodes, but the following key questions remain: 1) How to

*See our full paper (Chandak et al. 2022) for more details.
**Contributed to the paper equally
Copyright © 2022, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.
[1]Any node that is *not* required to be visited is referred to as a *Steiner node*. A path may choose to visit a Steiner node if it helps in either finding feasible solutions or reducing the cost of travel.

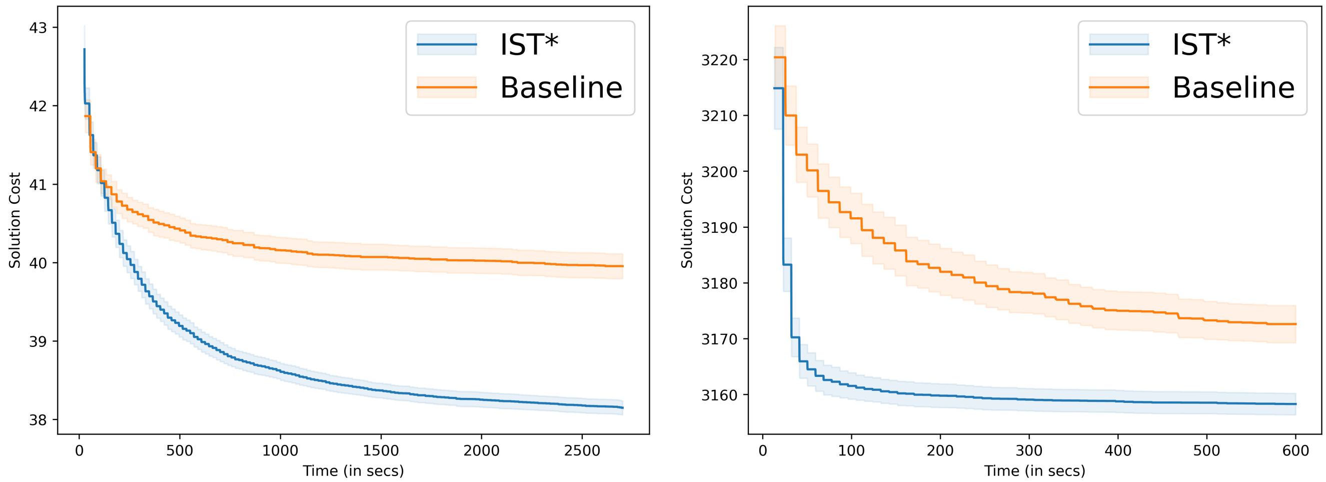

Figure 2: Experimental results (99% confidence interval) for two types of environments: a) $[0, 1]^8$ real-vector space (left); b) $SE(3)$ abstract environment (right).

sample the space so that the costs of the edges joining the nodes in T can be estimated quickly so that we can get a desired Steiner tree? 2) Should we estimate the cost of all the edges or can we ignore some edges and focus our effort on edges we think will likely end up in the Steiner tree?

Informed Steiner Trees (IST^*)

IST* iteratively alternates between sampling points in the search space and pruning edges. Throughout its execution, a Steiner tree is maintained which is initially empty but eventually spans the nodes in T, possibly including a subset of sampled points.

IST* relies on two key ideas. *First*, finding a Steiner tree spanning T commonly involves finding a Minimum Spanning Tree (MST) in the metric completion[2] of the nodes in T. The MST is at most the cost of the optimal MGPF path, and thus can be used to derive a 2-approximation guarantee. For any two distinct terminals u, v, we maintain a lower bound and an upper bound[3] on the cost of the edge (u, v). Using these bounds and cycle properties of an MST, we identify edges which will never be part of the MST (Pruning). This allows us to *only* sample regions corresponding to the edges that are likely to be part of the MST. We further bias our sampling by assigning a suitable probability distribution over the search space based on the bounds on the cost of the edges (See Fig. 1). *Second*, as the algorithm progresses, a new set of points are added to the search graph in each iteration. Each new sample added may facilitate a lower-cost feasible path between terminals requiring us to frequently update the Steiner tree. To address this efficiently, we develop an *incremental* version of the Steiner tree algorithm while maintaining its properties. Since this incremental approach correctly finds a Steiner tree, as the number of sampled points tends to infinity, IST* provides an asymptotic 2-approximation guarantee for MGPF.

We use the sampling procedure developed in *Informed RRT** (Gammell, Barfoot, and Srinivasa 2018) to choose points from selected regions in our approach. Informed sampling in synergy with pruning enables faster convergence to the optimal MST solution than uniform sampling.

Results

We evaluated the performance of IST^* against a baseline that involves densifying a roadmap using PRM* (Karaman and Frazzoli 2011) for a fixed amount of time via uniform random sampling, and then running S* (S*-BS variant was used) on the resulting roadmap to obtain an MST-based Steiner tree (Chour, Rathinam, and Ravi 2021). For a given environment, 50 terminals were randomly chosen in free space and both algorithms were repeated 50 times. It was ensured that the combined time spent in growing the roadmap and running S* exhausted a chosen time limit. The planners ($\mathtt{IST}^*$ and the Baseline) were implemented in Python 3.7 using OMPL (Şucan, Moll, and Kavraki 2012) v1.5.2 on a desktop computer running Ubuntu 20.04, with 32 GB of RAM and an Intel i7-8700k processor. Results for an experiment on a $[0, 1]^8$ real-vector space environment with a $[0.1, 0.8]^8$ central obstacle, and an $SE(3)$ environment are shown in Figure 2. We can observe in both subfigures, IST^* finds a cheaper Steiner tree in less time compared with the baseline.

Future Work

In this work, we introduced Informed Steiner Trees to find an asymptotic 2-approximation solution to the general case of the MGPF problem. We plan to explore obtaining effective lower bounds on the shortest path cost between two terminals instead of using the Euclidean distance.

[2]The metric completion here is a complete weighted graph on all the nodes in T where the cost of an edge between a pair of nodes in T is the minimum cost of a path between them.

[3]Upper bound is the cost of a feasible path from u to v.

Acknowledgments

This material is also based upon work partially supported by the National Science Foundation under Grant No. 2120219. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Best, G.; Faigl, J.; and Fitch, R. 2016. Multi-robot path planning for budgeted active perception with self-organising maps. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 3164–3171. IEEE.

Chandak, N.; Chour, K.; Rathinam, S.; and Ravi, R. 2022. Informed Steiner Trees: Sampling and Pruning for Multi-Goal Path Finding in High Dimensions.

Chour, K.; Rathinam, S.; and Ravi, R. 2021. S*: A Heuristic Information-Based Approximation Framework for Multi-Goal Path Finding. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 31, 85–93.

Edelkamp, S.; Secim, B. C.; and Plaku, E. 2017. Surface inspection via hitting sets and multi-goal motion planning. In *Annual Conference Towards Autonomous Robotic Systems*, 134–149. Springer.

Gammell, J. D.; Barfoot, T. D.; and Srinivasa, S. S. 2018. Informed sampling for asymptotically optimal path planning. *IEEE Transactions on Robotics*, 34(4): 966–984.

Gammell, J. D.; Srinivasa, S. S.; and Barfoot, T. D. 2014. Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2997–3004. IEEE.

Gammell, J. D.; Srinivasa, S. S.; and Barfoot, T. D. 2015. Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs. In *2015 IEEE international conference on robotics and automation (ICRA)*, 3067–3074. IEEE.

Janoš, J.; Vonásek, V.; and Pěnička, R. 2021. Multi-goal path planning using multiple random trees. *IEEE Robotics and Automation Letters*, 6(2): 4201–4208.

Karaman, S.; and Frazzoli, E. 2011. Sampling-based algorithms for optimal motion planning. *The international journal of robotics research*, 30(7): 846–894.

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

Kou, L.; Markowsky, G.; and Berman, L. 1981. A fast algorithm for Steiner trees. *Acta Informatica*, 15(2): 141–145.

Kuffner, J.; and LaValle, S. 2000. RRT-connect: An efficient approach to single-query path planning. In *Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No.00CH37065)*, volume 2, 995–1001 vol.2.

Macharet, D. G.; and Campos, M. F. 2018. A survey on routing problems and robotic systems. *Robotica*, 36(12): 1781–1803.

McMahon, J.; and Plaku, E. 2015. Autonomous underwater vehicle mine countermeasures mission planning via the physical traveling salesman problem. In *OCEANS 2015-MTS/IEEE Washington*, 1–5. IEEE.

Mehlhorn, K. 1988. A faster approximation algorithm for the Steiner problem in graphs. *Information Processing Letters*, 27(3): 125–128.

Otto, A.; Agatz, N.; Campbell, J.; Golden, B.; and Pesch, E. 2018. Optimization approaches for civil applications of unmanned aerial vehicles (UAVs) or aerial drones: A survey. *Networks*, 72(4): 411–458.

Şucan, I. A.; Moll, M.; and Kavraki, L. E. 2012. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4): 72–82. https://ompl.kavrakilab.org.