Risk Aware Graph Search With Uncertain Edge Costs

Ryan Skeele, Jen Jen Chung, Geoffrey A. Hollinger, and Kagan Tumer

***Abstract*— In this paper we present a novel approach to searching a graph with probabilistic edge costs. By incorporating uncertainty information into the graph search we perform risk-aware planning. We present the results through a simple search domain, and report the improved results compared to traditional single heuristic search techniques (*A**, *D**, and *greedy*).**

I. Introduction

Robots are becoming more integral to our everyday lives, but the transition into the unstructured world from the laboratory environment has proved challenging. With this transition there is a greater need for quick, reliable path planning methods, specifically under uncertainty. Planning under uncertainty allows for robustness in unstructured environments. We introduce a method, Risk-Aware Graph Search (RAGS), for finding paths through graphs with uncertain edge costs. Our method bridges the gap between traditional search methods and risk-aware planning.

Effectively searching through a graph with known edges has been extensively researched, and applies to many different applications. We aim to expand the graph search utility by allowing for uncertainty in the graph with risk-aware planning. Our novel approach searches over uncertain edge costs with known distributions, to find the best paths when there are no optimal path guarantees.

Traditional graph search methods (such as A* and D*) search over deterministic costs [1], [2] when traversing a graph. This approach ignores valuable information when dealing with uncertainty in edge costs. While there has been some work on risk-aware planning [3] [4], most work that does involve uncertainty is strictly concerned with the belief state of the robot's location [5] [6].

The main novelty of this paper is the introduction of a nonmyopic graph search algorithm for risk-aware planning. We present the search results of our method, compared to A*, D* and a greedy implementation. The results show that RAGS is more reliable than these existing methods.

II. Algorithm

RAGS leverages the confidence of edge costs to lower the probability of a poor path cost. There are two major steps to accomplish this, the first is to perform an initial search to find the set of non-dominated paths (Section II-A). Following this, we perform risk-aware planning during path execution as information of the neighboring edge costs are updated with the actual cost (Section II-B).

This work was supported by NASA grant NNX14AI10G.

The authors are with the School of Mechanical, Industrial and Manufacturing Engineering, Oregon State University, OR, 97330, USA skeeler@onid.oregonstate.edu, {jenjen.chung, geoff.hollinger, kagan.tumer}@oregonstate.edu

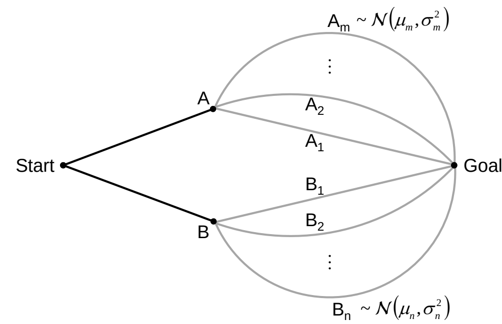

Fig. 1: Setup of one-to-one comparison for sequential look-ahead planning. Given the path cost distributions, we can directly compute the likelihood that traveling from Start to Goal through B will ultimately yield a cheaper path than traveling via A.

A. Bounding

Search algorithms like A* use total ordering of vertices in a graph based on cost to find a single optimal path. This is necessary because a point can have an infinite set of paths to it on a graph. Unfortunately we can't use the same technique because we are searching over two objectives (cost and variance). Instead we impose a partial ordering using Equation 1, which allows us to minimize the search space without pruning an optimal path. Equation 1 checks if a partial path is dominated. A path is dominated if both objectives are worse, as in Equation 1. If this is the case we no longer consider that path. This is a similar method as the one described in [5].

$$A < B \leftrightarrow (A.c < B.c) \wedge (A.\sigma^2 < B.\sigma^2) \quad (1)$$

Where A and B are partial paths, A.c is the cost (euclidean distance + additional cost) and A.σ^2 is the variance of path A cost.

B. Risk Comparison

Once the possible paths are generated, we must determine the best route to take. This means we need a way to compare each neighbor. We can compare by calculating the probability of a better path existing at that neighbor, which is calculated by integrating the paths associated with that neighbor.

Consider the setup in Fig. 1, the probability that a cheaper path exists from Start to Goal via vertex B over vertex A can be computed by considering the following: The probability

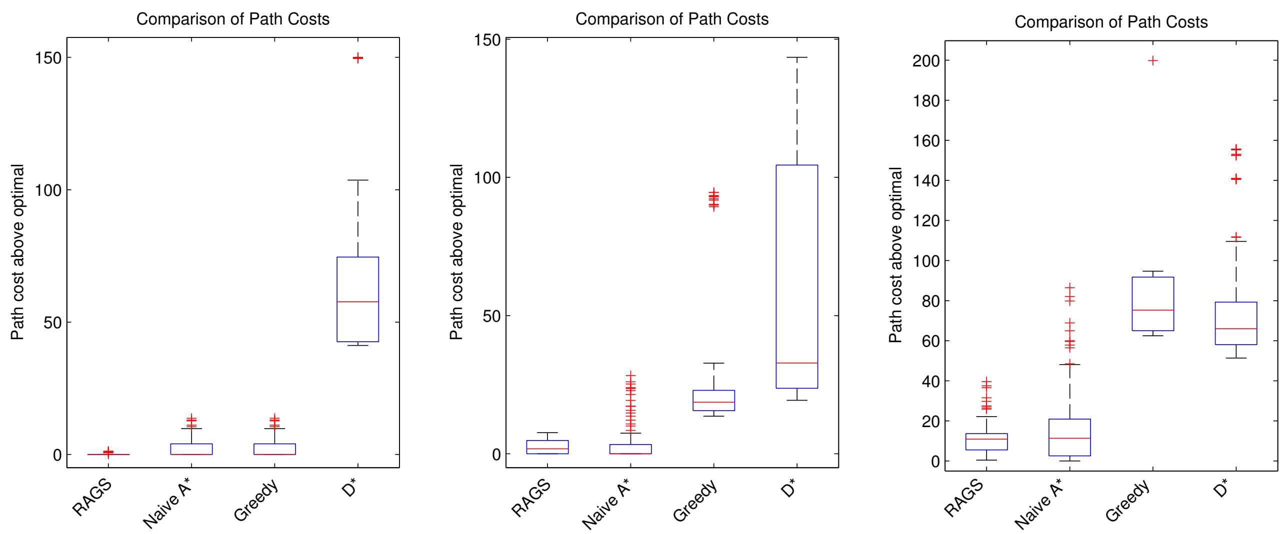

(a) (b) (c)

Fig. 2: Three plots showing search results over a PRM, with edge variances drawn uniformly between 0 and $\{5, 10, 20\}$ respectively.

that the best path from A to Goal has cost x is

$$\sum_{i=1}^{m} P\left(c_{A_i} = x;\, c_{A_j} > x, \forall j \neq i\right)$$
$$= \sum_{i=1}^{m} \left(\frac{1}{\sigma_{A_i}\sqrt{2\pi}} \exp\left(-d\left(A_i\right)^2\right) \prod_{j=1}^{m} \frac{1}{2}\mathrm{erfc}\left(d\left(A_j\right)\right) \right), \quad (2)$$

where c_{A_i} is the cost of the path from A to Goal through path i, and $d(\cdot) = \frac{x-\mu}{\sigma}$. The probability that at least one path from B to Goal has cost less than x is

$$1 - P\left(c_{B_i} > x, \forall i = \{1, \ldots, n\}\right)$$
$$= 1 - \prod_{i=1}^{n} \frac{1}{2}\mathrm{erfc}\left(d\left(B_i\right)\right). \quad (3)$$

If the cost of traveling from Start to A or B is c_{A_0} and c_{B_0}, respectively, then the probability that traveling from Start to B will yield a cheaper path to the Goal is the integral of the product of (2) and (3) over all possible values of x,

$$\int_{-\infty}^{\infty} \sum_{i=1}^{m} P\left(c_{A_i} = x;\, c_{A_j} > x, \forall j \neq i\right) \cdot$$
$$\left[1 - P\left(c_{B_i} > \left(x - \left(c_{B_0} - c_{A_0}\right)\right)\right), \forall i = \{1, \ldots, n\}\right] \mathrm{d}x. \quad (4)$$

III. Experiment Setup

The search algorithms were tested on a set of graphs generated with a uniform random distribution of 100 vertices over a space 100 x 100 in size, and connected according to the PRM* radius [7]. Edge costs were represented by normal distributions with mean equal to the Euclidean distance between vertices plus an additional cost drawn from a uniform random distribution over $[0, 100]$. The variance of each distribution was drawn from a uniform distribution over $\left[0, \sigma^2_{max}\right]$, where σ^2_{max}=$\{5, 10, 20\}$ for the three separate sets of experiments. Note that a minimum cost of the Euclidean distance was enforced in the following experiments. The start vertex was defined at $(0, 0)$, with the goal at $(100, 100)$.

We compared RAGS against a *naïve A** implementation, a *greedy* approach and *D**. During path execution the true costs of immediate neighboring edges become observable. Naïve A* finds and executes the lowest-cost path based on the mean edge costs and does not perform any replanning. The greedy search is performed over the set of non-dominated paths and selects the cheapest edge to traverse at each step, while D* replans over non-dominated paths at each step given the new edge cost information.

IV. Results

In Figure 2 we show the results for 100 trials, where each trial draws new edge costs from the same distribution. This distribution is generated for each plot as described in Section III. In these plots the variance of the edge costs increases (left to right) and we can see the overall trend of RAGS with a lower mean and variance in final path cost. Naïve A* also performs well but is more prone to outliers of more expensive paths (especially as the edge uncertainty increases), due to any edges with high variance along it's path. RAGS mitigates against skewed results like A*'s by choosing safer routes with lower variances, demonstrating the benefit of risk aware planning. Greedy does quite poorly because it does not look ahead and is fallible to traversing paths with high costs/variances and few path alternatives. D* also shows its fallibility to getting stuck along paths with high variance similar to A* and greedy.

V. Conclusion

In this paper we present a novel algorithm for searching through graphs with uncertain edge costs. The results of our work show that RAGS reduces the risk of a higher path cost, by doing risk-aware planning.

REFERENCES

[1] P. Hart, N. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *Systems Science and Cybernetics, IEEE Transactions on*, vol. 4, no. 2, pp. 100–107, July 1968.

[2] S. Koenig and M. Likhachev, "Fast replanning for navigation in unknown terrain," *Robotics, IEEE Transactions on*, vol. 21, no. 3, pp. 354–363, June 2005.

[3] L. Murphy and P. Newman, "Risky planning on probabilistic costmaps for path planning in outdoor environments," *Robotics, IEEE Transactions on*, vol. 29, no. 2, pp. 445–457, 2013.

[4] A. A. Pereira, J. Binney, G. A. Hollinger, and G. S. Sukhatme, "Risk-aware path planning for autonomous underwater vehicles using predictive ocean models," *Journal of Field Robotics*, vol. 30, no. 5, pp. 741–762, 2013.

[5] A. Bry and N. Roy, "Rapidly-exploring random belief trees for motion planning under uncertainty," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 723–730.

[6] H. Kurniawati, D. Hsu, and W. S. Lee, "SARSOP: Efficient point-based pomdp planning by approximating optimally reachable belief spaces." in *Robotics: Science and Systems*, 2008.

[7] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *The International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, 2011.