

Problem Space Analysis for Plan Library Generation and Algorithm Selection in Real-Time Systems

Robert Holder

University of Maryland Baltimore County

Abstract

Computing solutions to intractable planning problems is particularly problematic within real-time domains. One approach to this problem includes off-line computation of contingency plans. However, because complex domains preclude creating a comprehensive library, a system must choose a subset of all possible plans to include. Strategic selections will ensure that the library contains an appropriate plan for encountered situations.

This work discusses preliminary investigations into a scheme in which problem space analysis drives the creation of an efficient plan library. For complex problems, an exact analysis of the problem space is not feasible, and an efficient means of creating an approximate analysis is required. Thus, this work proposes the development of algorithms to efficiently generate and leverage the problem space analysis of complex planning problems.

Introduction

Many visitation planning problems can be mapped to a Dynamic Traveling Saleman Problem (DTSP), in which a system must plan a route to visit a set of potentially changing locations. When a new location becomes known, DTSP planners typically use heuristics to append the new locations to the previously computed route. Depending on the placement and quantity of these new locations, the accuracy of the approximate solution degrades. Instead of computing solutions at runtime, an alternate approach is to precompute solutions prior to runtime. One could imagine ideally precomputing a solution for every possible combination of potential new locations, but this is obviously impractical for large problems. However, it is upon this ideal that this approach is based.

This paper describes early experiments with creating a mapping between a sampling of problem instances and their solutions, with the intent of interpolating for the remainder of the problem space. These preliminary experiments demonstrate that a perfect library includes only a small percentage of all possible solutions. More importantly, experiments show that approximations of reasonable quality can be obtained from a small number of samples. The promise of these early results suggest that more sophisticated sampling techniques that leverage information acquired in previ-

ous samples, such as importance sampling and active learning, would lead to better results. Finally, this domain exhibits problem-specific hints can also assist with generating the problem space analysis.

Related Work

Minimizing the solutions required to achieve competent coverage of a problem space is well studied within Case-Based Reasoning (CBR) literature. Typically, a CBR system will encounter a problem and store the solution for future use. CBR is normally used in domains with discrete representations, although this is not always the case (Ram & Santamariá 1997). In most cases, CBR does not truly pre-plan; rather, all its solutions are generated during runtime.

Contingency planning is an alternative approach for generating plans for situations in which a plan may fail. One classic approach to contingency planning is Schopper's universal plans (Schoppers 1987) in which a solution to every possible situation is stored in a plan library. However, the potential drawback for this technique is the sheer number of states that must be considered. One alternative is to determine the necessary contingencies to plan for by calculating an expected *disutility* for an action that fails (Onder & Pollock 1996).

Problem Space Analysis

The problem space analysis currently consists of three maps: the Problem-Solution Map (PS Map), the Solution-Problem-Utility Map (SPU Map), and the Solution-Similarity Map (SS Map). For brevity, only the PS Map will be discussed in detail, and the other maps will be briefly mentioned.

An example of a PS Map is shown in Figure 1. This map shows the solutions for a set of 5-city DTSP problems. Four of the cities are fixed, as indicated by diamonds, and the axes represent the potential locations of the fifth city. The path starts in the middle at (0,0). For each of the possible locations of the fifth city (assuming integer coordinates), the shortest route is generated as the solution. Finally, each unique solution, consisting of a sequence of city identifiers, is assigned a color and plotted. For example, instances of the problem with the solution 0-1-5-3-2-4 might all be colored yellow.

This map assists with plan library creation by showing the minimum number of solutions required for optimal compe-

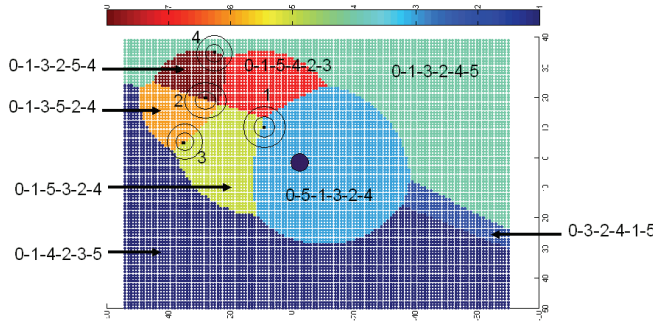


Figure 1: Problem-Solution Map of 5-city DTSP with marked city locations and solutions

tency across the problem space. We see that only eight solutions are required, representing just fewer than 7% of the 120 possible solutions. This is encouraging, but, for large problems, storing 7% of the possible solutions is not feasible.

The SPU map addresses this by suggesting regions in which several optimal solutions can be replaced with fewer suboptimal solutions while preserving reasonable plan quality. The SS Map attempts to further reduce the library size by discovering two or more solutions that are similar enough to be combined into one parameterized solution.

Preliminary Work & Results

Preliminary work using random sampling demonstrates that an accurate map can be generated with relatively few samples. In one experiment, the PS Map from figure 1 is sampled at various rates ranging from .0001 to .9, and solution interpolation is done by polling neighbors within a 15-unit radius. Finally, the percentage of correct solutions when compared to the original PS Map is plotted against the sample rate. Sampling at a .0001 rate yielded an accuracy rate of .544, while a .05 sample rate resulted in a .913 accuracy. The results of a follow up three-trial experiment focus on the sample rate between .0001 and .01 are shown in figure 2. Here we see that a .005 sample rate can generate a PS Map approximation that is a 80% match of the ideal PS Map.

Summary & Future Directions

This work presents preliminary work toward approximating the solution space for a complete array of problem instances using sampling and interpolation. These early experiments show that high solution accuracy can be obtained from a small sample of problem instance solutions. These experiments also suggest approaches for increasing the accuracy of the approximation, such as biasing samples towards areas of higher interest. Allocating additional samples near these typically heterogeneous regions seems a promising means of increasing the approximation accuracy. Also, replacing nearest neighbor classification with a more sophisticated approach will likely lead to better results. Finally, blending the

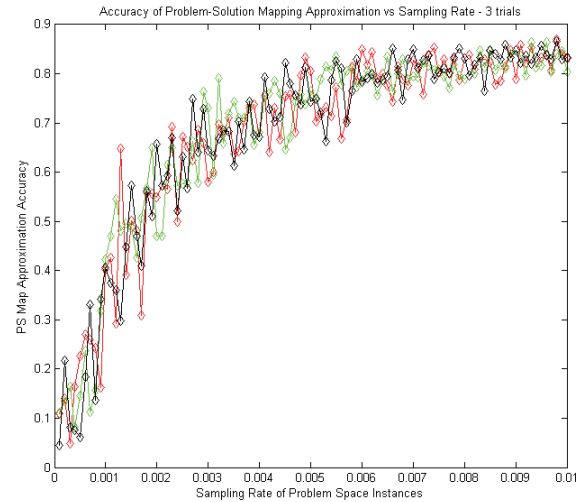


Figure 2: PS Map approximation accuracy at various sample rates (3 trials)

sampling and solution generation steps through active learning or importance sampling techniques is another consideration. It is currently unclear how well these techniques extend to the higher dimensions of more complex problems. In this case, finding an efficient means to accurately approximate the problem space analysis with a low sample rate becomes critical.

The application of these techniques to other domains is another consideration. For example, a wireless sensor network (WSN) must reconfigure itself to maintain a communications topology as targets move and sensor energy levels fluctuate. A predetermined library consisting of mappings from environment conditions to network configuration could reduce the amount of coordination necessary within the network, thus reducing communications needs and extending the life of the network. The interdependence of individual sensor policies may result in irregularities in the problem space, creating challenges when approximating the problem space analysis maps.

References

- Onder, N., and Pollack, M. E. 1996. Contingency selection in plan generation. In *Plan Execution: Problems and Issues: Papers from the 1996 AAAI Fall Symposium*, 102–108. AAAI Press, Menlo Park, California.
- Ram, A., and Santamaría, J. C. 1997. Continuous case-based reasoning. *Artif. Intell.* 90(1-2):25–77.
- Schoppers, M. J. 1987. Universal plans for reactive robots in unpredictable environments. In McDermott, J., ed., *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, 1039–1046. Milan, Italy: Morgan Kaufmann Publishers, Inc.: San Mateo, CA, USA.