Research Project Proposal
Multi-Robot Coverage of Modular Environments

Mirko Salaris, mirko.salaris@mail.polimi.it

1. Introduction to the problem

The Multi-Robot Coverage topic lies in the area of Multi-Robot Systems, which is a research area at the intersection of Artificial Intelligence and Autonomous Mobile Robotics. It studies methods and techniques to obtain advantages from the coordination and cooperation of multiple robots and to enable the execution of tasks that are inherently executable only by multiple agents.

In Multi-Robot Coverage the standard problem is finding optimal paths, w.r.t. time or distance metrics, that allow agents equipped with covering tools of finite size to completely cover a known environment. Formally, it is defined as the planning of a number of paths over the environment such that, when followed by the agents, all the points fall under the covering tool of at least one agent, the constraints of the environment, i.e., obstacles, are considered, and the solution is optimal w.r.t. a certain metric.

The coverage problem, and as a consequence the advantages brought by the multiple-robot approach, is driven by the practical need to physically pass over all the specified area (for example in applications like lawn mowing, floor cleaning or robotic de-mining [4]), to just gather data about the environment (e.g., water quality monitoring [14]), or to address search and rescue applications.

1.1. Preliminaries

Coverage can actually be solved by single robots, but Multi-Robot Coverage is an approach to coverage that exploits the usage of multiple robots to reduce the coverage time, to offer robustness (e.g., supporting the loss of a robot), and to enable coverage in special settings or under particular constraints [8]. The increased potentiality of the multi-robot approach comes with an increased algorithmic complexity due to methods that must take into account coordination and relations between different robots. Indeed, while there exists a polynomial-time coverage algorithm for the single robot case as shown by [16], finding a minimum travel time solution for the general multiple robot setting is an NP-hard problem [13].

We now need to introduce the multiple Traveling Salesmen Problem (mTSP), which is an extension of the more traditional Traveling Salesman Problem (TSP) to a setting with multiple salesmen instead of one. Here we present the usual definition of the mTSP [1]: given a set of nodes, let there be m agents located at a single initial node, called depot. The remaining nodes are called 'intermediate nodes'. The mTSP consists of finding tours for all the m agents, which all start and end at the depot, such that each intermediate node is visited exactly once and the total cost of visiting all nodes is minimized. The cost metric can be defined in terms of distance or time. In the area of Multi-Robot Systems we are usually interested in minimizing the total time of execution.

[10] provides an algorithm with a constant approximation factor of 2 for the Generalized, Multiple Depot, Multiple Travelling Salesman Problem (GMTSP), in the case of symmetric costs and with the triangle inequality satisfied. For the mTSP, [7] provides a tour-splitting heuristic that yields an approximation factor of $\frac{5}{2} - \frac{1}{m}$ relying on the $\frac{3}{2}$ TSP approximation by [5]. An open question is whether better approximation algorithms can be found.

A promising direction of research is to look for approximated algorithms that exploit features and characteristics of specific settings, for instance the presence of constraints on the structure of the environment. We want to investigate a setting that involves a modular environment with repeated identical sub-structures, which, we deem, allows for potential computational improvements (because of repetitions) and, as a consequence, for approximations with tighter bounds.

1.2. Problem description and motivations

We are looking for an efficient strategy to do coverage, with multiple robots, on a modular environment. 'Modular environment' means that the environment is characterized by repeated identical subparts that can be thought as modules, all equal to each other, of which the environment is made. This kind of structure is exemplified by that of a residential building which consists of multiple floors, all with the same floor plan, connected by one or more staircases. The different robots, equipped with a covering tool of finite size, are assumed to be all equivalent and capable of moving at constant speed through the whole environment, in such a way that we are able to abstract from the kinematic and dynamic constraints of a real scenario and to focus only on the algorithmic aspects. All the robots are initially placed at the same starting location. We define a solution as a set of paths, each of which corresponds to one robot, such that the first and the last points of each path coincide with the starting location of the robots and such that, when all the robots follow their paths, all the points of interest of the environment fall under the covering tool of at least one robot. We are interested in a solution with minimum completion time, where the coverage is considered finished when all the robots have navigated their paths.

2. Main related works

To the best of my knowledge, the problem of coverage in modular environments with repeated identical sub-structures has never been directly approached. Here, I present some works that are in some sense related and from which the problem arises.

An important theoretical result is that every mTSP can be approximately solved through a corresponding TSP formulation [2, 6, 12, 15]. The corresponding formulation is obtained creating m copies of the original depot, each connected to the other nodes exactly as the original depot, and manipulating the cost matrix. The TSP solution obtained on this new graph is forced to have m tours: the TSP path will go through each copy of the original depot m times. If we "cut" the TSP path every time it passes through a copy of the original depot, we obtain m paths, each of them starting from a copy of the depot and ending in a copy of the depot. These m paths are the mTSP (approximated) solution. We can already see how the issues in the mTSP are two: the partition of nodes among the different agents and the computation of the optimal paths. This is exactly how the problem is approached in the literature.

Indeed, a common approach for coverage problems in multi-agent settings is to group the nodes into m clusters, so that each cluster represents a set of adjacent nodes that can be visited by a single agent, whose path in the cluster can be optimized as in a standard TSP [3, 9, 11]. This idea of clustering is very used mainly because of the computational advantages: the search space of a routing problem with N nodes is a function of $N!$. A decomposition of the problem into k clusters means that the average number of nodes in a cluster is N/k. Therefore the search space for each cluster is $(N/k)!$ and the total search space is a function of $k \times (N/k)!$, which is much lower than $N!$ [3].

The work of [3] is very interesting for settings in which the completion time should be minimized, like in our case. In their work, the authors address the mTSP with the objective of balancing the workload among the agents and finding out the optimal number of agents required to cover the set of nodes. They first define the cluster length as the distance traveled by an agent to cover the cluster and return to the depot and then they estimate the optimal number of clusters k, with k inside a predefined range, taking into account a metric over the clusters lengths.

In our setting, the structure of the environment provides a hint on the partition of the nodes among the agents. Indeed, being the environment structured with repeated identical parts, we already have the equivalent of the clusters: each subpart can be considered as a cluster. So, with respect to what we find in the literature, the difference of our problem setting is that we already have the clusters and we don't have to build them. We are instead left with the problem of assigning the clusters to the agents.

3. Research plan

The purpose of the research is to investigate the potential improvements over the bounds of approximations given by path planning algorithms for multi-robot coverage problems enabled by considering environments characterized by repeated identical sub-structures, called modules.

The research is mainly of theoretical nature, concerning the study of algorithms, their complexity, and their guaranteed properties. We expect to conduct a numerical analysis in the last stages of the research.

The research will develop in three phases. In the first phase, we will look for a method that will provide a solution (i.e., an assignment of nodes to robots) for the case in which modules are considered as indivisible parts of the environment, meaning that each module will be assigned to one and only one robot. In this case, we expect to find an efficient algorithm able to assign each module to a single robot, with minimal overlap of routes among agents and that minimizes the overall completion time.
Later, we will address a relaxed version, in which each module can be shared between agents to further decrease the completion time, providing tighter bounds. The output of this phase will be an algorithm which we will refer to as 'without integrity constraints' to highlight the difference with the previous phase, where only an integer number of modules can be assigned to each robot.
In a third phase, we will consider, similarly to what is done in the literature, how to start from the solution of a global TSP over all the nodes and split it considering heuristics on the modules. Indeed, we expect the optimal solution to change with respect to the number of modules, how they are connected to each other, and their internal structure in terms of features such as the diameter of their graph representation or their internal optimal and approximated TSP.
After this three main phases, we expect to conduct a numerical analysis to complete the theoretical results.

Given the theoretical nature of the research we do not plan to use any experimental metrics, but we will evaluate our research output considering the approximation bounds guaranteed by our algorithms w.r.t. the optimal mTSP solution.

In Figure 1 we show the qualitative expected timeline of the research development. We expect the three phases to have a marginal amount of overlap and to last approximately 2 months each. The second phase is expected to last slightly longer not for an inherent difficulty but only for overlapping personal commitments of the author. We will start to write the paper before the end of the numerical analysis because most of the results should be already well assessed by that time.

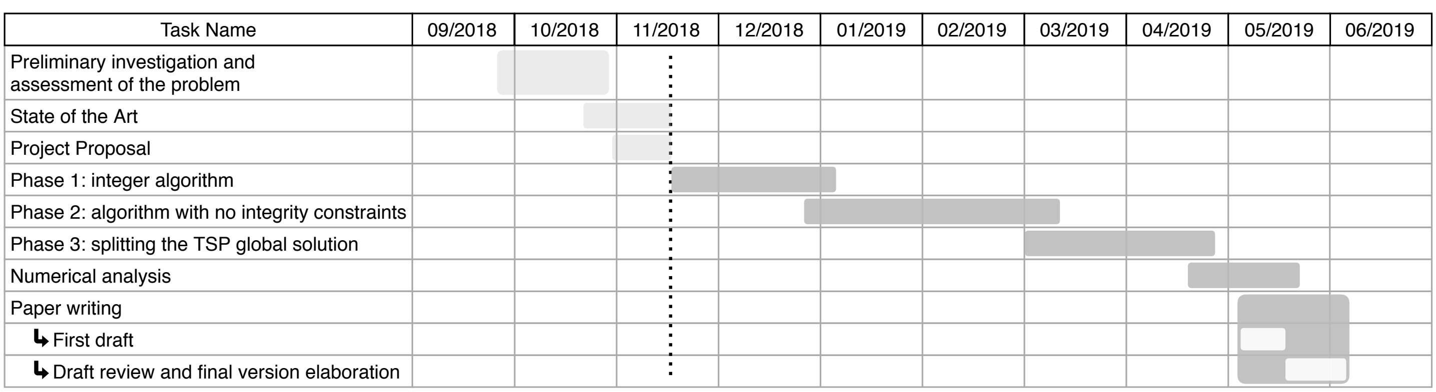

Figure 1: Gantt chart – expected timeline of the research development

References

[1] Bektas, T. The multiple traveling salesman problem: an overview of formulations and solution procedures. Omega 34, 3 (2006), 209–219.

[2] Bellmore, M., and Hong, S. Transformation of multisalesman problem to the standard traveling salesman problem. Journal of the ACM (JACM) 21, 3 (1974), 500–504.

[3] Chandran, N., Narendran, T., and Ganesh, K. A clustering approach to solve the multiple travelling salesmen problem. International Journal of Industrial and Systems Engineering 1, 3 (01 2006), 372–387.

[4] Choset, H. Coverage for robotics–a survey of recent results. Annals of mathematics and artificial intelligence 31, 1-4 (2001), 113–126.

[5] Christofides, N. Worst-case analysis of a new heuristic for the travelling salesman problem. Tech. rep., Carnegie-Mellon Univ Pittsburgh Pa Management Sciences Research Group, 1976.

[6] Eilon, S., Watson-Gandy, C. D. T., Christofides, N., and de Neufville, R. Distribution management-mathematical modelling and practical analysis. IEEE Transactions on Systems, Man, and Cybernetics SMC-4, 6 (Nov 1974), 589–589.

[7] Frederickson, G. N., Hecht, M. S., and Kim, C. E. Approximation algorithms for some routing problems. In Proceedings of the 17th Annual Symposium on Foundations of Computer Science (Washington, DC, USA, 1976), SFCS '76, IEEE Computer Society, pp. 216–227.

[8] Karapetyan, N., Benson, K., McKinney, C., Taslakian, P., and Rekleitis, I. Efficient multi-robot coverage of a known environment. In Proceedings of the 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (Sept 2017), pp. 1846–1852.

[9] Latah, M. Solving multiple TSP problem by k-means and crossover based modified ACO algorithm. International Journal of Engineering Research and Technology 5 (02 2016), 430–434.

[10] Malik, W., Rathinam, S., and Darbha, S. An approximation algorithm for a symmetric generalized multiple depot, multiple travelling salesman problem. Operations Research Letters 35, 6 (2007), 747–753.

[11] Nallusamy, R., Duraiswamy, K., Dhanalaksmi, R., and Parthiban, P. Optimization of non-linear multiple traveling salesman problem using k-means clustering, shrink wrap algorithm and meta-heuristics. International Journal of Nonlinear Science 9, 2 (2010), 171–177.

[12] Orloff, C. Routing a fleet of m vehicles to/from a central facility. Networks 4, 2 (1974), 147–162.

[13] Rekleitis, I., New, A. P., Rankin, E. S., and Choset, H. Efficient boustrophedon multi-robot coverage: an algorithmic approach. Annals of Mathematics and Artificial Intelligence 52, 2-4 (2008), 109–142.

[14] Singh, A., Krause, A., Guestrin, C., and Kaiser, W. J. Efficient informative sensing using multiple robots. Journal of Artificial Intelligence Research 34 (2009), 707–755.

[15] Svestka, J. A., and Huckfeldt, V. E. Computational experience with an m-salesman traveling salesman algorithm. Management Science 19, 7 (1973), 790–799.

[16] Xu, A., Viriyasuthee, C., and Rekleitis, I. Efficient complete coverage of a known arbitrary environment with applications to aerial operations. Autonomous Robots 36, 4 (2014), 365–381.