

Constrained Exploration for Studies in Multirobot Coordination

Nidhi Kalra, Dave Ferguson, and Anthony Stentz

The Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania
{nkalra,dif,tony}@cmu.edu

Abstract—We are interested in the challenges of long-term, long-distance tight multirobot coordination. In this paper, we discuss the Constrained Exploration domain that pushes the boundaries of tightly-coordinated multirobot teams. We formalize this problem and consider the role of intentional coordination approaches in addressing its numerous challenges. Finally, we present three intentional coordination approaches to solving different aspects of the domain.

I. INTRODUCTION

Many multirobot domains require that teammates closely coordinate their actions throughout execution. In the existing literature, this tight coordination has almost exclusively been achieved using reactive or behavior-based techniques where coordination is a byproduct of robots following simple rules [1]. These approaches are fast, simple, and robust, and they work well in domains such as object transport where teammates can respond to each other in simple, prescribed ways and when planned coordination is not necessary.

We are interested in the challenges of long-term, long-distance tight coordination that does require planning and that perhaps cannot be solved by these methods. In this paper, we discuss the Constrained Exploration domain that pushes the boundaries of tightly-coordinated multirobot teams. We formalize this problem and consider the role of intentional coordination approaches in addressing its numerous challenges. Finally, we present three intentional approaches to solving different aspects of the domain.

II. THE CONSTRAINED EXPLORATION DOMAIN

Consider a team of robots tasked with exploring a hazardous environment. For robustness, information dissemination and retrieval, or tasking, we require that robots always remain within communication access to some base station in the environment, either directly or by using teammates as relay nodes. In essence, the team must maintain an ad-hoc network, which may require that individual robots tightly-coordinate to ensure continuous connectivity for all members. We call this domain *Constrained Exploration* (CE).

We formalize this domain as a special case of the Multi-depot Traveling Salesman Problem (MD-TSP) in continuous space. The MD-TSP is defined by a set S of salesmen that start at different depots, a set of cities C , and some cost for traveling between cities. Each city must be visited by one salesman, and

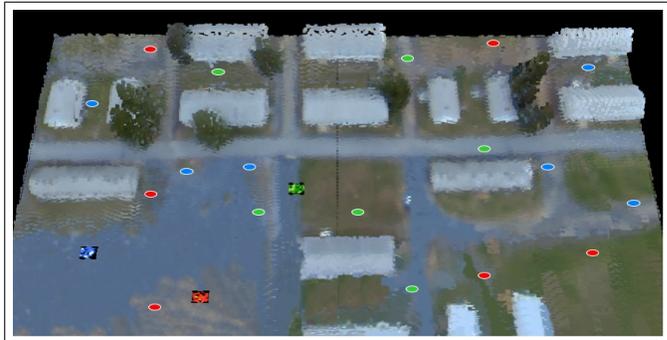


Fig. 1. An example of Constrained Exploration with three robots and eight cities marked as discs. The left-most robot carries the communication antenna, so every team member must be able to communicate with it at all times. In this instance, communication requires a clear line of sight.

the total distance traveled by the salesmen must be minimized. To solve the MD-TSP, the cities must be allocated to the salesmen and each salesman must construct a tour of minimum cost. MD-TSP has been shown to be \mathcal{NP} -hard [1].

For Constrained Exploration, the salesmen S correspond to members of the robot team R and the cities correspond to exploration sites. We then add a set A of communication sites and two functions that describe the cost of communication at any point during execution: $\text{cost-to-commsite}(r_i, a_j)$ is the cost of communication between robot r_i and communication site a_j and $\text{cost-to-robot}(r_i, r_k)$ is the cost of communication between robots r_i and r_k . Typically, these cost functions will depend on the relative locations of the robots and sites, which may change over time. We then require that at all times, each robot's cost of communicating with at least one site remain below some threshold \mathcal{T}_c . The cost of communication for robot r_i at any time is defined as:

$$\text{cost}(r_i) = \min \left[\begin{array}{l} \min_{a_j \in A} \text{cost-to-commsite}(r_i, a_j), \\ \min_{r_k \in R} \text{cost-to-robot}(r_i, r_k) + \text{cost}(r_k) \end{array} \right]$$

Defined in this way, Constrained Exploration closely couples task decomposition, allocation, and execution. Thus, it is difficult to solve one part of this problem well without simultaneously solving the other parts. Constrained Exploration can also be mapped to a number of other domains such

as communication-constrained reconnaissance and cleanup, visibility-constrained pursuit evasion and perimeter sweeping, and distance-constrained transportation.

Constrained Exploration presents several challenges to multirobot coordination. Firstly, as mentioned, task allocation schemes are required that consider the continuous interactions between robots during execution when making assignments. Thus, determining which robot should do a task depends on both when a candidate robot will do that task and what its teammates will be doing at that time. Secondly, in realistic instances, robots will be operating in uncertain and dynamic environments. Consequently, tightly-coordinating robots will have to intelligently share information: if too little is shared, they may fail to react to conditions that threaten their success, but if too much is shared, they may be bogged down trying to process it. There is also scope for planning algorithms that repair the existing tightly-coupled solutions when new information arrives. Thirdly, malfunction and failure are inescapable facts of every real world domain; in tightly-coordinated teams in particular, mission success depends on the simultaneous success of multiple teammates. In this domain, robots may malfunction and leave teammates unexpectedly without communication access. To be robust, the team must be able to detect failures, propagate this information to all the affected team members, and then restructure the team to compensate.

III. EXAMPLE APPROACHES

The Constrained Exploration domain is essentially a multi-agent planning problem in which the actions of robots are tightly coupled. Thus, approaches that consider the long-term impacts of teammates' decisions (i.e. that plan) are likely to fare better than those approaches that behave myopically. However, optimally solving the general Constrained Exploration problem as defined above is currently infeasible even for small teams. Nevertheless, it is possible to solve different parts of the problem, albeit approximately, using current algorithms. In the remainder of this paper, we present three example approaches that address a simplification of the problem: we assume away the task allocation aspects and focus on mechanisms to coordinate the tours.

A. Centralized Approaches

Centralized approaches to multirobot problems plan for the entire team, simultaneously taking into account the environment and the interactions of all team members at all times. Thus, in theory, they can produce optimal solutions. In reality, centrally coordinating more than a few robots becomes intractable since the complexity of algorithms is usually exponential in the number of robots. Centralized approaches also tend to be brittle to failure and respond slowly to change. However, distributed approaches have been developed that opportunistically use centralized planning to improve solutions [2], [3]. Consequently, we believe that centralized planners can play an important role in the success of distributed approaches.

To this end, we have investigated two different approaches to making centralized planning for Constrained Exploration

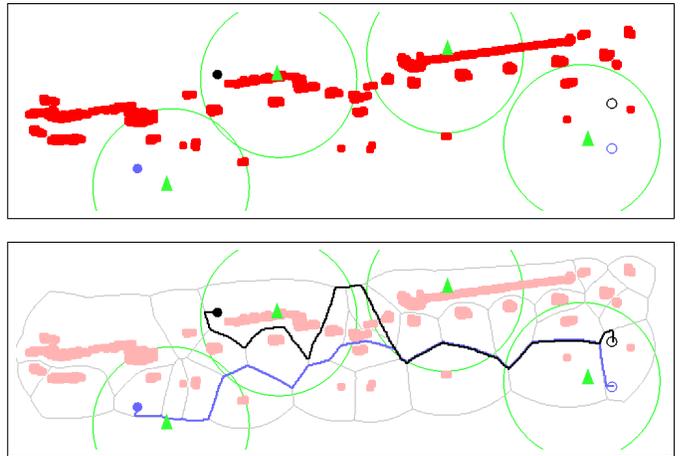


Fig. 2. Using a GVD to reduce the search space in a Constrained Exploration domain. (Top) The two robots are represented by filled circles and must move to their goal city which are open circles, while always remaining in contact with at least one communication tower. The four communication towers are marked as green triangles and their range is indicated by the thin circles around them. (Bottom) The search space was reduced to just the GVD of the environment, in gray. The robots' paths that solve this instance are overlaid.

more tractable. The first method is to reduce the search space to a roadmap; for this, we use the Generalized Voronoi Diagram (GVD) [4]. A GVD is a road-map that provides all possible path homotopies in an environment containing obstacle regions. We use a GVD to facilitate centralized planning using A* for a typical Constrained Exploration problem shown in Figure 2. In this cluttered 650×200 environment, two robots were each pre-assigned one city which they had to reach while keeping in contact with one of four communication sites. Here, the cost-to-commsite and cost-to-robot functions returned 0 if the robots were in range of the sites and each other, respectively, and infinity otherwise. \mathcal{T}_c was set to 1. The size of the full state space was 1.7×10^{10} ; the GVD reduced the state space to 1.3×10^6 . Thus, whereas A* had taken 94 seconds on the full search space, it required only 0.36 seconds on the reduced search space.

A second option for dealing with the large search space is to use probabilistic planners such as Rapidly-exploring Random Trees (RRTs) [5]. RRTs combine random sampling of the search space with biased sampling around the goal to grow trees that quickly plan through vast, high-dimensional spaces. For example, Fig. 3 shows an RRT and the corresponding set of paths for 3 robots moving through a 300×600 environment containing obstacles through which communication cannot occur. Thus, cost-to-commsite and cost-to-robot functions return 0 if there is line of sight and infinity otherwise. As before, \mathcal{T}_c is set to 1. In this instance of the problem, there is exactly one communication site that is mobile and carried by one of the robots on the team. In effect, this requires that the team remain fully connected at all times.

The size of the state space for this problem was 6×10^{15} , however the RRT was able to find a path in under a second¹.

¹Using a 1.5 GHz Powerbook G4.

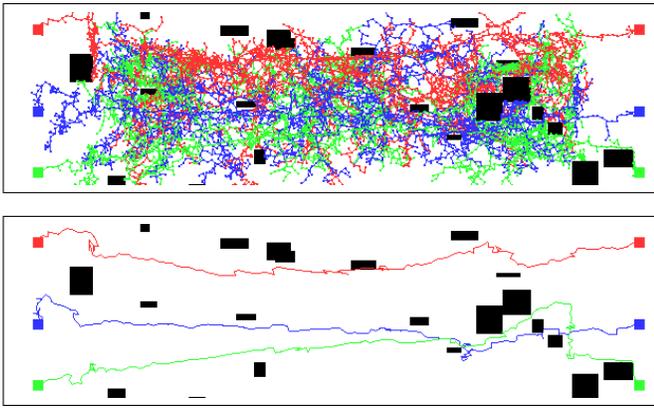


Fig. 3. Constrained Exploration. A path is planned for three robots from one side to the other of an environment, while maintaining line-of-sight communication constraints among the team. (top) The RRT. To show the 6D RRT in only two dimensions, we split into three different-colored 2D trees. Each 2D tree encodes all the information about a single robot from the 6D tree, including edges between parent and child nodes. The black regions depict areas through which communication cannot be made. (bottom) The corresponding solution path for each robot (before smoothing).

In fact, even for large teams, RRTs are able to generate paths in a reasonable amount of time: to plan for 10 robots in a 10^{54} state space, the RRT generated a path in 13 seconds.

B. Distributed Approaches

Distributed approaches are popular in current multirobot research. In these approaches, each robot acts as an independent agent and makes decisions with local information about its state and its environment. Consequently, they tend to be more robust to failure, more flexible, and more tractable; however, solutions may also be highly suboptimal. For example, Wagner and Arkin’s behavior-based approach to this problem [6] is simple and flexible, but does not employ planning which is necessary for complex instances. To harness the benefits of centralized approaches in distributed systems, market-based approaches have been designed to centrally plan over small subsets of the team when time and resources permit [2], [3]. In market-based frameworks, robots model an economy of self-interested individuals that buy and sell tasks and resources to maximize personal profit [1]. This redistribution of tasks and resources simultaneously results in lower cost solutions for the team.

We have applied the Hoplites framework developed by Kalra et. al. [3] to Constrained Exploration. Hoplites is a market-based framework used for tightly-coupled multirobot planning problems in which robots buy and sell centrally-planned solutions to smaller parts of the problem space. Kalra showed that this approach significantly outperformed behavior-based approaches in a related domain. Figure 1 illustrates an instance of Constrained Exploration on which we used Hoplites. A set of eighteen cities has been pre-allocated to the team of three robots, indicated by cities and robots of matching color. As in the previous instance, the communication antenna is carried by a robot and line of sight is required for communication. Figure 4 is a snapshot from a simulation showing how robots

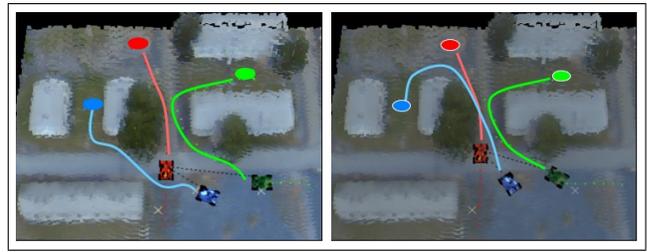


Fig. 4. A snapshot illustrating the Hoplites framework applied to a Constrained Exploration problem. (Left) With the current paths, the red/left and green/right robots will lose communication with the blue/center robot. (Right) Having detected this problem, both the blue and the red robots propose solutions. The blue robot’s solution (not shown) requires that the red robot change its course to the other side of the obstacle. This is a high cost solution because it leaves the green robot stranded. The red robot’s solution requires that the blue robot change its course. This is a lower cost solution and is successfully sold by the red robot to the blue robot.

sell small centralized plans to facilitate tight coordination. We have demonstrated Hoplites in simulation for Constrained Exploration using three robots and up to twenty cities.

IV. FUTURE DIRECTIONS

We have only just scratched the surface of the many research avenues possible in the many variations of the Constrained Exploration domain. Our future work includes extending both the GVD and RRT-based centralized approaches to handle efficient replanning, and also incorporating task allocation into the Hoplites approach. Although these are only a few possibilities, we hope they will motivate others to explore the challenges of this domain. We look forward to the resulting research.

V. ACKNOWLEDGMENTS

This work was sponsored by the U.S. Army Research Laboratory, under contract “Robotics Collaborative Technology Alliance” (contract number DAAD19-01-2-0012). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies or endorsements of the U.S. Government. Dave Ferguson is supported in part by a National Science Foundation Graduate Research Fellowship.

REFERENCES

- [1] M. B. Dias, N. Kalra, R. Zlot, and A. Stentz, “Market-based multirobot coordination: A survey and analysis,” *IEEE Special Issue on Multirobot Systems*, 2006 (forthcoming).
- [2] M. B. Dias, “Traderbots: A new paradigm for robust and efficient multirobot coordination in dynamic environments,” Ph.D. dissertation, Robotics Institute, Carnegie Mellon University, January 2004.
- [3] N. Kalra, D. Ferguson, and A. Stentz, “Hoplites: A market-based framework for complex tight coordination in multi-robot teams,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2005.
- [4] H. Choset and J. Burdick, “Sensor based planning, part I: The generalized voronoi graph,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 1995.
- [5] S. LaValle and J. Kuffner, “Randomized kinodynamic planning,” *IJRR*, vol. 20, no. 5, pp. 378–400, 2001.
- [6] A. Wagner and R. Arkin, “Multi-robot communication-sensitive reconnaissance,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, April 2004.