Optimization Approaches to Multi-robot Planning and Scheduling

Kyle E. C. Booth
Department of Mechanical & Industrial Engineering
University of Toronto, Toronto, Ontario, Canada
kbooth@mie.utoronto.ca

The study and use of multi-robot teams has become more prevalent within academia and industry as the capability and autonomy of these systems continues to improve (Arai, Pagello, and Parker 2002). With high levels of progress already made concerning the control of individual robots, the acknowledged advantages of multi-robot systems (MRS) has resulted in considerable research attention in the last couple of decades in efforts to build more efficient systems for their coordination (Gerkey and Matarić 2004). These advantages include resolving more complex tasks, increasing the speed at which tasks can be completed, and enhancing the level of system reliability and redundancy present within single-robot solutions (Yan, Jouandeau, and Cherif 2013).

This thesis, though in early stages of development, concerns the field of study within MRS known as *multi-robot task allocation* (MRTA) (Gerkey and Matarić 2004). Specifically, it proposes to investigate the integration of techniques from the optimization literature, namely mixed-integer and constraint programming, within architectures for MRTA. This research area aims to solve multi-robot coordination problems pertaining to task distribution to robot resources and the temporal scheduling of tasks on such resources. These problems have a wide variety of real-world applications including planetary exploration (Mataric and Sukhatme 2001), airport and harbor transhipment (Alami et al. 1998), and emergency response (Østergård, Matarić, and Sukhatme 2001).

Multi-robot Task Allocation

Given a team of cooperating robots, a set of tasks that need to be completed, and a problem-specific cost function, the most fundamental instance of MRTA involves determining a mapping of tasks to robots such that the cost function is minimized. Indeed, when the number of robots and tasks are equal and the mapping is one-to-one, the problem can be represented by the classical *linear assignment problem*, solvable in $O(n^3)$ time with a modificiation of the Hungarian method (Kuhn 1955). However, MRTA problems often contain more complex objective functions or the need for tasks to be allocated *and* scheduled on available resources. In such cases, the underlying problem frequently becomes provably $\mathcal{NP}$-Hard, taking on the form of other classical problems within the combinatorial optimization literature such as the Multiple Traveling Salesman Problem (m-TSP) (Papadimitriou 1977). Moreover, the presence of complex constraints (e.g. precedence relationships) between tasks within and across robot schedules further contributes to the difficulty of solving such problems. These *time-extended* allocation problems with *task-dependencies* (Korsah, Stentz, and Dias 2013) are often approached with sophisticated heuristic techniques that sacrifice solution optimality in favour of faster convergence.

Existing Approaches

Early efforts to develop solutions for MRTA include iterated assignment architectures such as ALLIANCE (Parker 1998) and M+ (Botelho and Alami 1999). These architectures employ dispatch-style algorithms where single tasks are assigned and executed before subsequent allocations are made. Various decentralized and fully-distributed techniques have also been proposed for these problems, notably the market-based (Dias et al. 2006) and auction-based (Gerkey and Matari 2002) methods developed within the robotics community. More recently, there have been efforts to use linear and mixed-integer programming (MIP) techniques from the operations research (OR) community to solve MRTA problems (Korsah et al. 2012), largely due to attractive bounds on solution quality, though these methods have not been fully exploited as of yet. Constraint programming (CP) has been proposed as a suitable candidate approach for these problems (Van Hentenryck and Saraswat 1996), however, the application of CP to multi-robot task planning and scheduling has been, to the best of our knowledge, limited.

Research Focus

The first component of our ongoing research focuses on the development and application of optimization-based methods to solve single and multi-robot task planning problems. We investigate the use of MIP and CP techniques to produce high-quality robot task plans that yield provable bounds on solution quality. The suitability of these 'model-and-solve' techniques is computationally assessed, and methods for improving algorithm performance through specialized modeling (e.g., symmetry breaking constraints, auxiliary variables) and search manipulation (e.g., branching rules, variable and value ordering heuristics) are implemented. As part of our future research we plan to develop more specialized algorithms based on hybrid approaches, incorporating concepts from the OR, CP, and AI communities to further enhance performance. We also aim to study the use of our optimization-based methods for plan repair and replanning, algorithm extensions that look to address real-world uncertainty associated with MRTA applications.

The second component of our research explores the integration of these optimization-based techniques within real-world robot architectures, combining environmental perception, achieved through on-board sensors, with our task planning system to achieve truly autonomous decision-making. Within this integration, the task planning methods developed contribute as a high-level mission planner for the system, allocating tasks to each robot team member as well as specifying when and where such tasks should be executed. These integration efforts are realized using the open source Robot Operating System (ROS) (Quigley et al. 2009), which enables the effective management of communication among individual robot subsystems or between multiple robots. In this architecture, we use peer-to-peer communication graph renderings where nodes represent individual subsystems or robots and arcs represent connections between them. Using ROS-specific implementation details such as *messages* and *topics*, we are able to effectively facilitate system communication. Communication is used to coordinate individual robot function, inform the system of changes within the environment, as well as deliver task allocations and commands to the various robot resources. We validate the utility of our methods using both simulated environments and testing on physical robots. For simulated experimentation, we use ROS Visualization software to model the robots and high-level task planner. These experiments are primarily aimed at improving the performance of our task planning methods. Physical robot implementation validates the real-world function of our task planner as well as important lower-level systems including autonomous navigation, path-planning, and object detection. For these tests we use OpenSlam's (openslam.org) GMapping to create our environment map via simultaneous localization and mapping (SLAM).

Research efforts thus far have shown promise through our work on a single-robot task scheduling problem involving a socially-assistive robot facilitating human-robot interactions (HRI) within a retirement home (Booth et al. 2016). The problem involves reasoning about disjoint time windows, robot travel times, intra-schedule task dependencies, and robot energy levels. For the problem studied, CP has been shown to be the dominant technology, finding high-quality solutions in much shorter runtimes than both MIP and temporal planning techniques. This optimization-based approach is integrated on the social robot *Tangy*, using the ROS architecture as previously described. The integration is tested experimentally on a number of realistic scenarios, demonstrating its physical utility as a viable robot task planning alternative. As a natural progression to our work, research efforts are underway to extend this work to multi-robot variants of the problem.

Conclusion

We explore the use of optimization-based methods for multi-robot task allocation (MRTA) problems. Our contributions and ongoing research consist of two components: i) algorithmic development, and ii) system integration and testing.

For the first component, we investigate the use of 'off-the-shelf' applications of mixed-integer programming (MIP) and constraint programming (CP), as well as look into enhancing these approaches through specialized modeling, search manipulation, and hybridization. In the second component, we work towards integrating our planning and scheduling algorithms within a functioning robot architecture, using the Robot Operating System (ROS) to facilitate system communication. We test our methods using simulated environments, achieved with ROS Vizulization software, and validate their utility with implementation on physical robot systems. In future research, we plan to explore algorithmic extensions that incorporate replanning events and plan repair in efforts to address the uncertainty associated with real-world problems.

References

Alami, R.; Fleury, S.; Herrb, M.; Ingrand, F.; and Robert, F. 1998. Multi-robot cooperation in the martha project. *Robotics & Automation Magazine, IEEE* 5(1):36–47.

Arai, T.; Pagello, E.; and Parker, L. E. 2002. Editorial: Advances in multi-robot systems. *IEEE Transactions on robotics and automation* 18(5):655–661.

Booth, K. E.; Tran, T. T.; Nejat, G.; and Beck, J. C. 2016. Mixed-integer and constraint programming techniques for mobile robot task planning. *Robotics and Automation Letters, IEEE* 1(1):500–507.

Botelho, S. C., and Alami, R. 1999. M+: a scheme for multi-robot cooperation through negotiated task allocation and achievement. In *Robotics and Automation, 1999. Proceedings. 1999 IEEE International Conference on*, volume 2, 1234–1239. IEEE.

Dias, M. B.; Zlot, R.; Kalra, N.; and Stentz, A. 2006. Market-based multirobot coordination: A survey and analysis. *Proceedings of the IEEE* 94(7):1257–1270.

Gerkey, B. P., and Matari, M. J. 2002. Sold!: Auction methods for multirobot coordination. *Robotics and Automation, IEEE Transactions on* 18(5):758–768.

Gerkey, B. P., and Matarić, M. J. 2004. A formal analysis and taxonomy of task allocation in multi-robot systems. *The International Journal of Robotics Research* 23(9):939–954.

Korsah, G. A.; Kannan, B.; Browning, B.; Stentz, A.; and Dias, M. B. 2012. xbots: An approach to generating and executing optimal multi-robot plans with cross-schedule dependencies. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, 115–122. IEEE.

Korsah, G. A.; Stentz, A.; and Dias, M. B. 2013. A comprehensive taxonomy for multi-robot task allocation. *The International Journal of Robotics Research* 32(12):1495–1512.

Kuhn, H. W. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly* 2(1-2):83–97.

Mataric, M. J., and Sukhatme, G. S. 2001. Task-allocation and coordination of multiple robots for planetary exploration. In *In Proceedings of the 10th International Conference on Advanced Robotics*. Citeseer.

Østergård, E. H.; Matarić, M. J.; and Sukhatme, G. S. 2001. Distributed multi-robot task allocation for emergency handling. In *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, volume 2, 821–826. IEEE.

Papadimitriou, C. H. 1977. The euclidean travelling salesman problem is np-complete. *Theoretical Computer Science* 4(3):237–244.

Parker, L. E. 1998. Alliance: An architecture for fault tolerant multirobot cooperation. *Robotics and Automation, IEEE Transactions on* 14(2):220–240.

Quigley, M.; Conley, K.; Gerkey, B.; Faust, J.; Foote, T.; Leibs, J.; Wheeler, R.; and Ng, A. Y. 2009. Ros: an open-source robot operating system. In *ICRA workshop on open source software*, volume 3, 5.

Van Hentenryck, P., and Saraswat, V. 1996. Strategic directions in constraint programming. *ACM Computing Surveys (CSUR)* 28(4):701–726.

Yan, Z.; Jouandeau, N.; and Cherif, A. A. 2013. A survey and analysis of multi-robot coordination. *International Journal of Advanced Robotic Systems* 10.