DART: Diversity-enhanced Autonomy in Robot Teams

Nora Ayanian

Abstract This paper defines the research area of Diversity-enhanced Autonomy in Robot Teams (DART), a novel paradigm for the creation and design of policies for multi-robot coordination. While current approaches to multi-robot coordination have been successful in structured, well understood environments, they have not been successful in unstructured, uncertain environments, such as disaster response. The reason for this is not due to limitations in robot hardware, which has advanced significantly in the past decade, but in how multi-robot problems are solved. Even with significant advances in the field of multi-robot systems, the same problem-solving paradigm has remained: assumptions are made to simplify the problem, and a solution is optimized for those assumptions and deployed to the entire team. This results in brittle solutions that prove incapable if the original assumptions are invalidated. This paper introduces a new multi-robot problem-solving paradigm which relies on a diverse set of control policies that work together synergistically to make multi-robot systems more resilient in unstructured and uncertain environments.

1 Introduction

The field of multi-robot systems (MRS) is growing at a rapid pace. Research in MRS spans many different areas, including automated delivery [1–3], surveillance [4], and disaster response [5, 6]. There have also been many successful demonstrations of increasing numbers of robots [7–9, 11–13]. MRS have also been successfully deployed in the field including in warehousing [15], manufacturing [16], and entertainment [17]. While these outcomes show the promise of MRS, the environments in which MRS have been successful are highly controlled, and some are highly instrumented, enabling precise tuning of controllers and nearly perfect knowledge of environmental conditions.

Nora Ayanian
University of Southern California, Los Angeles, CA, USA e-mail: ayanian@usc.edu

Many environments where MRS could be beneficial are not highly controlled or equipped with the extensive infrastructure often necessary to coordinate large teams of robots with state-of-the-art algorithms. For example, containing wildfires, searching collapsed buildings, patrolling borders, monitoring infrastructure, and containing oil spills all occur in highly dynamic and unique environments (no two collapsed buildings are the same), with high uncertainty and little control over other non-robot agents in the environment. One of the most desirable benefits of MRS is *robustness*, wherein robots can compensate for loss of capabilities by relying on other robots in the team. However, the uncertainty of many real-world environments renders current algorithms, even those designed for robustness, ineffectual. *The reason for this is not due to limitations in robot hardware, but in how multi-robot problems are solved.* Many controllers are so specialized and optimized for specific capabilities and conditions that they cannot cope with uncertainty. Thus, the true benefits of robustness in teams of robots have yet to be achieved.

2 Motivation

In disaster response alone, the potential impact of autonomous MRS is substantial: 60,000 people die each year in natural disasters [18]. This makes robots an ideal tool for disaster response. In fact, DJI announced that *one* properly equipped drone can find a missing person more than *five times faster* than traditional search methods [19]. However, most robots used in search and rescue today are teleoperated [20], requiring trained operators which may not be nearby. Disaster response that is autonomous, without the need for an expert operator, can reduce response time and save more lives, especially when a trained operator may be hours away.

The potential applications of autonomous MRS go well beyond disaster response, including military, agriculture, transportation, manufacturing, and fulfillment applications. However, current solutions for MRS have not successfully transitioned from controlled environments such as laboratories or warehouse facilities to the inherently high uncertainty in these complex environments. Without infrastructure that provides communication and localization, and without knowledge of or control over the environment, current state-of-the-art methods fail.

While the field of MRS has advanced significantly, the same problem-solving paradigm has remained. First, the problem is defined. Next, complexity is reduced by making several assumptions to simplify the problem, such as terrain and communication range. Finally, an optimal solution to that specific problem is designed and applied to all the robots in the team. This paradigm (Fig. 1a) limits the capability of MRS to cope with real-world environments. The solutions are brittle, as the assumptions made are easily invalidated and the optimized controller is not designed for real environments. In the best case, the controller is able to overcome these challenges, but it is not a good solution to the problem. In the worst case, the controller cannot cope, causing mission failure, loss of high-value assets, and casualties; after all, if the same failed controller is applied to all robots, all of them will fail.

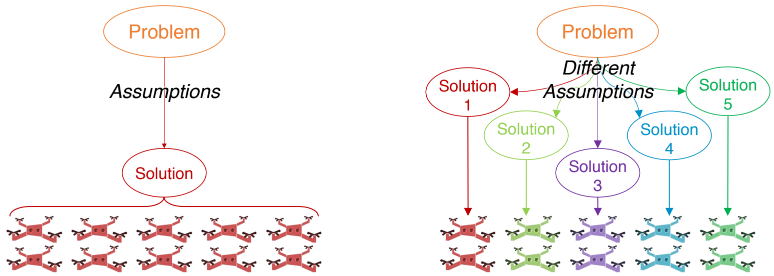

(a) Curent MRS Paradigm (b) Proposed MRS Paradigm

Fig. 1: (a) The current MRS problem-solving paradigm is linear, applying the same solution to all robots. (b) The proposed novel paradigm takes advantage of diversity in controllers to handle various scenarios

3 A Potential Solution

Instead of applying the same controller to all robots, a new approach leveraging diversity in policies within the robot team can allow MRS to better cope with uncertain environments. Using an ensemble of diverse control policies to accomplish a coordinated task within a single team of robots can enable the team to adjust to different conditions. For example, with two unmanned aerial vehicles (UAV) on a large security task, a natural result of using an ensemble of controllers is for one UAV to position itself high, to view the entire area, while the other UAV takes a closer look at areas of interest.

Diversity is well established as a way to improve the performance of human workgroups: studies have shown repeatedly that diverse groups outperform homogeneous groups [21–24]. Thus, *the current problem-solving paradigm in MRS does not reflect an effective approach to working in groups*. Instead of the current paradigm of solving problems and uniformly applying the solution to all robots as in Fig. 1a, several solutions to the problem under different assumptions and different styles of interaction should be developed and the best approaches combined to take advantage of their strengths under different conditions, as in Fig. 1b.

4 Current State of the Art

While diversity of robots with different physical embodiment or capabilities has previously been studied [32–34, 36], there has been relatively little exploration into diversity in control policies within a single team of robots. Most research in this area is a result of studying ants that take different roles in foraging and house hunting [37, 39, 40] or collective transport [41], and applied to similar problems in robotics. Unfortunately, in trying to model ant algorithms closely, these works do not take advantage of robot capabilities, including communication, sensing, and computation.

In Tang and Parker's ASyMTRe, robots take different roles depending on environmental conditions [42, 43], but the robots are all programmed to react the same. This leaves them vulnerable to unforeseen changes in capabilities or the environment, and does not enable robots to individually adjust their approaches.

A majority of work exploring control diversity in robots exists in behavior-based systems, most notably Balch's work in learning behavioral specialization for robot teams [44, 45]. Goldberg and Matarić evaluate multi-robot controllers based on the amount of interference and describe caste arbitration, where all robots have the same capabilities, but have different conditions for activating behaviors [46].

More recently, evolutionary robotics and agent-based systems have been appearing as a method for encouraging behavioral diversity and plasticity (individuals changing roles over time). Mouret and Doncieux review and benchmark published approaches to behavioral diversity, and show that fostering behavioral diversity substantially improves the evolutionary process in the investigated experiments, regardless of task [47]. Pugh *et al.* review quality diversity algorithms, which have resulted in a new class of algorithms that return an archive of diverse, high-quality behaviors in a single run [48]. Vassiliades and Christodoulou design behaviorally plastic agents (capable of switching between different behaviors in response to environmental changes) [49] and Umedachi *et al.* attempt to understand the underlying mechanism of the behavioral diversity of animals, then use the findings to build truly adaptive robots [50]. However, all of these approaches focus on training agents to act independently in the environment, and thus are not directly applicable to multi-robot problems where task completion relies on tight coordination, such as box-pushing, shape formation, wildfire containment, cooperative transport, etc. Furthermore, agents are trained in the environments where they will be used, which, especially in natural disasters, may not be possible.

Heterogeneity has also been studied extensively in insect and animal behavior. Jandt *et al.* study personality at multiple levels with regard to behavioral syndromes and insect societies, discussing fitness consequences of intra-colony behavioral variation [51]. Specifically, under varying environmental conditions, maintaining a mixture of individuals with different behavioral types may be more effective than individuals switching between behavioral types, which might be costly and inefficient. Slower, more accurate individuals can bring large quantities of food back to the colony when good abundance is constant, whereas faster "sloppier" individuals might be more efficient at exploiting resources in more frequently changing environments [53]. Burns and Dyer [55] found that ant colonies that maintain a mixture of different foraging types within a group allows colonies to respond more quickly to environmental fluctuation. On the other hand, maintaining a mixture of inflexible behavioral types can incur costs to the colony, such as overly aggressive types being aggressive to their own nestmates [58].

These results in insect and animal behavior studies point strongly to behaviorally heterogeneous teams having higher fitness in uncertain and dynamic environments, which has inspired many multi-robot approaches. However, there is a need to further the use of diversity as a tool for MRS, especially in tightly-coordinated tasks.

5 Some Open Problems in Diversity-enhanced Autonomy for Robot Teams

Much as human workgroups, as well as insects and animals, benefit from diversity in composition of the group, such variation of behavior would be beneficial for teams of robots operating in uncertain and unstructured environments. There exist many open problems in DART; some of the challenging open problems that must be addressed by the community are described here.

1. **Learning from Humans:** Humans provide a pool of diverse resources that can be tapped to develop diverse controllers that work well together. However, due to differences in human and robot capabilities (communication, locomotion, sensing, etc.), it is difficult to learn controllers by observing human in-person interaction. By limiting interaction to an interface (such as a mobile phone, tablet, or laptop), communication, locomotion, and sensing can be restricted to robot-like capabilities [59]. A major benefit of human-inspired controllers is the ability to communicate with and easily motivate study participants, as opposed to animal-inspired controllers. However, learning from human cooperation requires multi-agent learning tools for many agents. This is an area that is not yet well represented in the literature, save for several works [**?**, **?**, **?**].
2. **Deep Learning for MRS:** In order to learn from humans, or to learn directly from simulations, new machine learning tools must be developed for multi-agent systems. While some solutions exist in multi-agent learning, some focus on tasks that can be learned and completed alone [60, 61], and those that are suitable for tight coordination for a few (2-3) agents, are intractable for large numbers of agents that must tightly coordinate [62–64]. Tight coordination between a large team of agents, for example in wildfire containment, currently presents a significant computational challenge for existing multi-agent learning tools.
3. **Measures of Diversity and Fitness:** Taking inspiration from the study of behavioral diversity in social insect colonies, there is a need for understanding the impact of behavioral diversity on MRS in tightly coordinated tasks. To that end, measures of diversity and fitness must be developed that apply to MRS, such as Balch's Hierarchic Social Entropy [45]. Such tools will likely be task-specific at first, while the science of diversity-enhanced autonomy is established.
4. **Adjusting Policies Online:** To successfully utilize a diverse set of controllers, the team of robots must collectively reason about the role that each team member plays and automatically adjust their own roles to achieve an appropriately diverse team with an effective skill set. To do so, they must have the ability to measure the success of individual agents on a coordinated task, learning from their own and others' shortcomings and successes.

6 Discussion

This paper proposes a new research thrust that represents a paradigm-shift in problem-solving for multi-robot systems from a linear paradigm, where policies are optimized for a specific set of assumptions and applied to the entire team, to one where policies are developed with multiple sets of assumptions and exist synergistically within a team of robots. Such diversity in control policies will better prepare the team of robots for challenging environments, much like diversity in the knowledge base in human workgroups leads to higher quality solutions. Adoption of this new paradigm may lead to expanded success of multi-robot systems in the field, especially in unstructured and uncertain environments.

A small sample of open problems were discussed, but there exist many open problems in this space. By explicitly defining Diversity-enhanced Autonomy for Robot Teams, we hope to inspire the development of new tools for coping with uncertain, unstructured environments such as first response, precision agriculture, surveillance, and others.

Acknowledgements

This work was supported by NSF CAREER (IIS-1553726) and the Okawa Foundation Research Award. Special thanks to Gaurav Sukhatme, M. Ani Hsieh, and Fei Sha for conversations and guidance in formulating this line of research.

References

1. Ali-akbar Agha-mohammadi, N. K. Ure, J. P. How, and J. Vian. Health aware stochastic planning for persistent package delivery missions using quadrotors. In *IEEE/RSJ Int Conf Intelligent Robots and Systems*, pages 3389–3396, Sept 2014.
2. Kat Lonsdorf. Hungry? call your neighborhood delivery robot. NPR Morning Edition, `http://www.npr.org/sections/alltechconsidered/2017/03/23/520848983/hungry-call-your-neighborhood-delivery-robot`, 2017.
3. C. Sung, N. Ayanian, and D. Rus. Improving the performance of multi-robot systems by task switching. In *IEEE International Conf Robotics and Automation*, pages 2999–3006, May 2013.
4. April Glaser. These surveillance robots will work together to chase down suspects. Recode, `https://www.recode.net/2017/4/18/15264908/surveillance-robots-network-cornell-suspects`, 2017.
5. N. Schurr, J. Marecki, M. Tambe, P. Scerri, N. Kasinadhuni, and J. Lewis. The future of disaster response: Humans working with multiagent teams using defacto. In *AAAI Spring Symp. AI Technologies for Homeland Security*, 2005.
6. J. S. Jennings, G. Whelan, and W. F. Evans. Cooperative search and rescue with a team of mobile robots. In *Advanced Robotics, 1997. ICAR '97. Proceedings., 8th International Conference on*, pages 193–200, Jul 1997.

7. Timothy H. Chung, Michael R. Clement, Michael A. Day, Kevin D. Jones, Duane Davis, and Marianna Jones. Live-fly, large-scale field experimentation for large numbers of fixed-wing uavs. In *IEEE International Conference on Robotics and Automation*, pages 1255–1262, 2016.
8. April Glaser. Intel invented a way for a single operator to fly hundreds of drones at once. Recode, `https://www.recode.net/2016/11/4/13517550/intel-single-operator-fly-hundreds-drones-shooting-star`, 2016.
9. Sabine Hauert, Severin Leven, Maja Varga, Fabio Ruini, Angelo Cangelosi, Jean-Christophe Zufferey, and Dario Floreano. Reynolds flocking in reality with fixed-wing robots: Communication range vs. maximum turning rate. In *IEEE/RSJ Intl Conf Intelligent Robots and Systems*, pages 5015–5020. IEEE, 2011.
10. Intel and Ars Electronica Futurelab. Drone 100. World record for the most UAVs airborne simultaneously, `http://www.spaxels.at/show_drone100.html`, 2015.
11. Alex Kushleyev, Daniel Mellinger, Caitlin Powers, and Vijay Kumar. Towards a swarm of agile micro quadrotors. *Autonomous Robots*, 35(4):287–300, 2013.
12. J. A. Preiss, W. Hoenig, G. S. Sukhatme, and N. Ayanian. Crazyswarm: A large nano-quadcopter swarm. In *IEEE Intl Conf Robotics and Automation*, 2017.
13. M. Rubenstein, C. Ahler, and R. Nagpal. Kilobot: A low cost scalable robot system for collective behaviors. In *IEEE International Conference on Robotics and Automation, ICRA 2012, 14-18 May, 2012, St. Paul, Minnesota, USA*, pages 3293–3298, May 2012.
14. Gábor Vásárhelyi, Csaba Virágh, Gergo Somorjai, Norbert Tarcai, Tamás Szörényi, Tamás Nepusz, and Tamás Vicsek. Outdoor flocking and formation flight with autonomous aerial robots. In *IEEE/RSJ Int Conf Intelligent Robots and Systems*, pages 3866–3873, 2014.
15. R. D'Andrea and P. Wurman. Future challenges of coordinating hundreds of autonomous vehicles in distribution facilities. In *IEEE International Conference on Technologies for Practical Robot Applications*, pages 80–83, Nov 2008.
16. James R. Hagerty. Meet the new generation of robots for manufacturing. *Wall Street Journal*, June 2, 2015.
17. Brian Barret. Disney's latest attraction? 300 drones flying in formation. Wired, `https://www.wired.com/2016/11/disneys-latest-attraction-300-drones-flying-formation/`, 2016.
18. Charles Kenny. *Why Do People Die In Earthquakes? The Costs, Benefits And Institutions Of Disaster Risk Reduction In Developing Countries*. The World Bank, 2009.
19. DJI. DJI documents faster search and rescue responses with drones. DJI Newsroom, `http://www.dji.com/newsroom/news/dji-documents-faster-search-and-rescue-responses-with-drones`, 2016.
20. Yugang Liu and Goldie Nejat. Robotic urban search and rescue: A survey from the control perspective. *J. Intell. Robotics Syst.*, 72(2):147–165, November 2013.
21. L. R. Hoffman. The group problem-solving process. In L. Berkowitz, editor, *Group Processes*, number 101-114. Academic Press, New York, 1978.
22. L.R. Hoffman and N.R.F. Maier. Quality and acceptance of problem solutions by members of homogeneous and heterogeneous groups. *Journal of Abnormal and Social Psychology*, 62:401–407, 1961.
23. Charlan Nemeth. Differential contributions of majority and minority influence. *Psychological Review*, 93:23–32, 1986.
24. Susan Jackson. Team composition in organizations. In S. Worchel, W. Wood, and J. Simpson, editors, *Group Process and Productiv-ity*. Sage, London, 1992.
25. Ivan D. Steiner. *Group Process and Productivity*. Academic Press, San Diego, 1972.
26. Charles A. O'Reilly and Sylvia Flatt. Executive team demography: Organizational innovation and firm performance. *Working paper, School of Business, University of California at Berkeley*, 1989.
27. Deborah Ancona and David Caldwell. Demography and design: Predictors of new product team performance. *Organization Science*, 3:321–341, 1992.
28. Karen A. Jehn. A multimethod examination of the benefits and detriments of intragroup conflict. *Administrative Science Quarterly*, 40(2):256–282, 1995.

29. Evert van de Vliert and Carsten K.W. de Dreu. Optimizing performance by conflict stimulation. *International Journal of Conflict Management*, 5(3):211 – 222, 1994.
30. Frank R. C. de Wit, Lindred L. Greer, and Karen A. Jehn. The paradox of intragroup conflict: A meta-analysis. *Journal of Applied Psychology*, 97(2):360–390, Mar 2012.
31. E. Diller, S. Floyd, C. Pawashe, and M. Sitti. Control of multiple heterogeneous magnetic microrobots in two dimensions on nonspecialized surfaces. *IEEE Transactions on Robotics*, 28(1):172–182, Feb 2012.
32. L. E. Parker. Alliance: an architecture for fault tolerant multirobot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, Apr 1998.
33. L. C. A. Pimenta, V. Kumar, R. C. Mesquita, and G. A. S. Pereira. Sensing and coverage for a network of heterogeneous robots. In *IEEE Conf Decision and Control*, pages 3947–3952, Dec 2008.
34. Jiangyang Huang, S. M. Farritor, A. Qadi, and S. Goddard. Localization and follow-the-leader control of a heterogeneous group of mobile robots. *IEEE/ASME Transactions on Mechatronics*, 11(2):205–215, April 2006.
35. B. Grocholsky, A. Makarenko, and H. Durrant-Whyte. Information-theoretic coordinated control of multiple sensor platforms. In *IEEE Intl Conf Robotics and Automation*, volume 1, pages 1521–1526 vol.1, Sept 2003.
36. A. Prorok, M. A. Hsieh, and V. Kumar. The impact of diversity on optimal control policies for heterogeneous robot swarms. *IEEE Transactions on Robotics*, 33(2):346–358, April 2017.
37. M. Dorigo, M. Birattari, and T. Stutzle. Ant colony optimization. *IEEE Computational Intelligence Magazine*, 1(4):28–39, Nov 2006.
38. M. Dorigo, G. D. Caro, and L. M. Gambardella. Ant algorithms for discrete optimization. *Artificial Life*, 5(2):137–172, April 1999.
39. K. Sugawara, T. Kazama, and T. Watanabe. Foraging behavior of interacting robots with virtual pheromone. In *IEEE/RSJ Int Conf Intelligent Robots and Systems*, volume 3, pages 3074–3079 vol.3, Sept 2004.
40. S. Berman, A. Halasz, V. Kumar, and S. Pratt. Bio-inspired group behaviors for the deployment of a swarm of robots to multiple destinations. In *Proceedings 2007 IEEE International Conference on Robotics and Automation*, pages 2318–2323, April 2007.
41. Ganesh P. Kumar, Aurélie Buffin, Theodore P. Pavlic, Stephen C. Pratt, and Spring M. Berman. A stochastic hybrid system model of collective transport in the desert ant aphaenogaster cockerelli. In *Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control*, HSCC '13, pages 119–124, New York, NY, USA, 2013. ACM.
42. Fang Tang and L.E. Parker. A complete methodology for generating multi-robot task solutions using asymtre-d and market-based task allocation. In *Proc. IEEE Conf. Robotics and Automation*, pages 3351–3358, Apr. 2007.
43. L. E. Parker and Fang Tang. Building multirobot coalitions through automated task solution synthesis. *Proceedings of the IEEE*, 94(7):1289–1305, July 2006.
44. Tucker Balch. Learning roles: Behavioral diversity in robot teams. In *AAAI Workshop on Multiagent Learning*, 1997.
45. Tucker Balch. Hierarchic social entropy: An information theoretic measure of robot group diversity. *Autonomous Robots*, 8(3):209–238, 2000.
46. Dani Goldberg and Maja J Mataric. Interference as a tool for designing and evaluating multirobot controllers. In *Proceedings AAAI*, pages 637–642, Providence, Rhode Island, 1997.
47. J. B. Mouret and S. Doncieux. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary Computation*, 20(1):91–133, 2016/11/03 2011.
48. Justin K. Pugh, Lisa B. Soros, and Kenneth O. Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.
49. Vassilis Vassiliades and Chris Christodoulou. Behavioral plasticity through the modulation of switch neurons. *Neural Networks*, 74:35 – 51, 2016.
50. Takuya Umedachi, Kentaro Ito, and Akio Ishiguro. Soft-bodied amoeba-inspired robot that switches between qualitatively different behaviors with decentralized stiffness control. *Adaptive Behavior*, 23(2):97–108, 2015.

51. Jennifer M. Jandt, Sarah Bengston, Noa Pinter-Wollman, Jonathan N. Pruitt, Nigel E. Raine, Anna Dornhaus, and Andrew Sih. Behavioural syndromes and social insects: personality at multiple levels. *Biological Reviews*, 89:48–67, 2014.
52. N. E. Raine and L. Chittka. The correlation of learning speed and natural foraging success in bumble-bees. *Proceedings of the Royal Society B: Biological Sciences*, 275:803–808, 2008.
53. L. Chittka, P. Skorupski, and N. E. Raine. Speed-accuracy tradeoffs in animal decision making. *Trends in Ecology and Evolution*, 24:400–407, 2009.
54. N. E. Raine and L. Chittka. No trade-off between learning speed and associative flexibility in bumblebees: a reversal learning test with multiple colonies. *PLoS One*, 7, 2012.
55. J.G. Burns and A.G. Dyer. Diversity of speed-accuracy strategies benefits social insects. *Current Biology*, 18:R953–R954, 2008.
56. A. P. Modlmeier and S. Foitzik. Productivity increases with variation in aggression among group members in temnothorax ants. *Behavioral Ecology*, 22:1026–1032, 2011.
57. J. N. Pruitt and S. E. Riechert. How within-group behavioural variation and task efficiency enhance fitness in a social group. *Proceedings of the Royal Society B: Biological Sciences*, 278:1209–1215, 2011.
58. M. W. J. Crosland. Variation in ant aggression and kin discrimination ability within and between colonies. *Journal of Insect Behavior*, 3:359–379, 1990.
59. A. Tavakoli, H. Nalbandian, and N. Ayanian. Crowdsourced coordination through online games (Late Breaking Report). In *ACM/IEEE Intl Conf Human-Robot Interaction*, Christchurch, New Zealand, Mar 2016.
60. Maja J Mataric. Learning to behave socially. In *Third international conference on simulation of adaptive behavior*, volume 617, pages 453–462, 1994.
61. Thomas Recchia, Jae Chung, and Kishore Pochiraju. Improving learning in robot teams through personality assignment. *Biologically Inspired Cognitive Architectures*, 3:51–63, 2013.
62. Olivier Buffet, Alain Dutech, and François Charpillet. Shaping multi-agent systems with gradient reinforcement learning. *Autonomous Agents and Multi-Agent Systems*, 15(2):197–220, 2007.
63. Christopher Amato, George D Konidaris, and Leslie P Kaelbling. Planning with macro-actions in decentralized pomdps. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, pages 1273–1280. International Foundation for Autonomous Agents and Multiagent Systems, 2014.
64. Christopher Amato, George Konidaris, Gabriel Cruz, Christopher A Maynor, Jonathan P How, and Leslie P Kaelbling. Planning for decentralized control of multiple robots under uncertainty. In *Robotics and Automation (ICRA), 2015 IEEE International Conference on*, pages 1241–1248. IEEE, 2015.