Active Recruitment Mechanisms for Heterogeneous Robot Teams in Dangerous Environments

Geoff Nagy and John Anderson(✉)

Autonomous Agents Laboratory, University of Manitoba,
Winnipeg, MB R3T 2N2, Canada
{geoffn,andersj}@cs.umanitoba.ca

Abstract. Using teams of autonomous, heterogeneous robots to operate in dangerous environments means increased cost-effectiveness and the ability to spread skills among team members. The high risk of loss in these domains is a challenge to team management. Teams must be able to recruit the help of other robots in the environment, while balancing searching with performing immediately useful work. This paper describes additions to a framework for dynamic team management in dangerous domains in order to support various levels of active search for useful agents while balancing useful work in the domain.

1 Introduction

Heterogeneity is useful in multi-robot teams because it may not be cost-effective to provide every team member with the most expensive equipment, and doing so would only increase the robots' design and control complexity [1]. Losses in dangerous environments can limit a team's skill set, and since many or even all robots are likely to eventually be damaged or lost in such environments, it is sensible and necessary to release replacement robots periodically. To accomplish this, teams must also be prepared to integrate new robots (possibly including previously lost units) and decide how much effort to expend into searching for them. More active approaches involving physical searches will yield faster results at the expense of immediately useful work. Passive approaches that rely on chance encounters will result in more useful work being performed by individual robots, but will lessen the chances of successfully locating desired robots [8]. Strategies in between these extremes would contain elements of both, such as simply calling for help wirelessly without physically searching. We have made additions to an existing framework [8] enabling teams of heterogeneous robots to use various recruitment strategies, and this paper focuses only on a specific subset of these. We examine in simulation how our framework performs in a dangerous domain: Urban Search and Rescue (USAR), the search for human victims inside a partly collapsed structure, under conditions of partial communication failure and robot loss.

R. Khoury and C. Drummond (Eds.): Canadian AI 2016, LNAI 9673, pp. 276–281, 2016.
DOI: 10.1007/978-3-319-34111-8_34

2 Previous Work

Existing work does not employ a wide range of recruitment strategies and many works assume that robots are always in communication range of each other [2,3,11], or focus on domains [4,5,12,14] that do not reflect real-world conditions that our work takes into account (e.g., communication failures or robot loss).

Krieger et al. [10] modelled foraging tasks using recruitment among homogeneous robots, but did not consider factors such as multiple robot or task types, or unreliable communication. Pitonakova et al. [13] endeavoured to discover under what conditions recruitment is beneficial for more complex foraging but did not explore elements of robotic heterogeneity or specialization despite the presence of different resource types. Kiener and Von Stryk [9] presented a task allocation framework for teams of highly heterogeneous and specialized robots, demonstrated with a small range of tasks. It contained no mechanisms to support the addition of new robots or the possibility of larger roles.

Gunn and Anderson [8] developed a framework allowing teams of heterogeneous robots to operate in complex simulated USAR environments, including unreliable communication and possible loss of robots. Teams of robots were assigned tasks by a leader to explore the environment and locate human victims. Losses were countered by the ability to rebalance roles (including leadership) upon losing or encountering robots. We have made additions to this framework allowing robots to search for useful skills as needed (as opposed to relying on chance encounters) and balance this activity with performing useful work. The next section describes relevant portions of the original framework and follows with the new mechanism we have added.

3 Methodology

3.1 Framework

In order to describe our recruitment mechanisms, the main components of our existing framework for dynamic team management and task allocation must be described. Space limits the background that can be presented here; the interested reader is directed to [8].

Robots. Our framework is designed to support heterogeneous robots that will not be equally suitable for all types of work. In our work, robots are able to determine how well-suited they are for a particular job and every robot maintains a priority queue of outstanding jobs. Robots can be assigned tasks by a team leader and can also discover them on their own. Newly discovered tasks can be executed by the robot itself or can be passed to a team leader for reassignment elsewhere if the robot is ill-equipped or too busy.

Tasks, Roles, Team Leaders, and Task Allocation. We define a *task* as a single piece of work that can be completed by a robot. Every robot in our work fills a *role*, which includes a description of the types of tasks associated with it.

As the milieu of team members changes due to losses or robot encounters, robots use their own (limited and possibly inconsistent) team knowledge and may switch roles in order to fill gaps in team abilities [8]. A leader robot, although not always reachable (and thus without the most up-to-date knowledge of the team), brings global perspective to a team to the degree that communication is available. The leader should be the most computationally-capable unit on a team, and in the original framework, is responsible for task assignments, updating known victim locations, and keeping a limited model of the current team.

The team leader uses roles as a heuristic to indicate default assignments to team members, but communication limitations (such as a delayed report of a role change) and robot losses can cause this to fail. Task broadcast is used as a fall-back, and members respond with their ability to complete the task irrespective of their role [7]. A robot receiving a request to complete a task may not always be able to oblige due to limited capabilities or an overfull task queue. While the original framework supported task assignments only from the team leader, our recruitment extensions add a mechanism whereby other team members unable to accomplish a task can recruit others to complete it.

3.2 Recruitment Strategies

Recruitment strategies vary in terms of how much effort robots expend towards locating others. Here we describe two mechanisms that represent points on a spectrum of more active recruitment mechanisms.

Concurrent recruitment increases the chances of encountering other robots while also completing useful work. Robots will perform tasks as normal while simultaneously broadcasting messages asking for help with other tasks. Other robots in wireless range, even if on another team, will accept the task if they are able, or offer to recruit for the task if they are not. A recruiter will examine these responses and attempt to assign a task to the most suitable robot who is able to complete the task themselves. If no such robot is available, the recruiter will attempt to offload the task to one of the other responding robots for recruitment. If no responses are received or the assignment fails (due to robot failures or unreliable communication) the recruiter continues to broadcast the request.

With *active recruitment*, a robot will perform a physical search for a robot that can complete a specific task, rather than completing pending tasks of its own. Physically searching for another robot has the advantage of providing increased opportunities for encountering others, at the cost of spending time away from useful work. In situations where a robot is given one task too many, or has been assigned a task for which it is poorly-equipped, it will enqueue the task and mark it to be completed via recruitment. This ensures that the recruitment task is weighed against other pending tasks. When recruiting, the robot performs a physical search to increase the chances of encountering another robot, and broadcasts a request for help and the task details. As with concurrent recruitment, the recruiter assigns the task to the most suitable robot who responds to the message. Once the task has been successfully assigned, the recruiter returns to its team's last known location and resumes normal tasks.

4 Experimental Evaluation and Design

We evaluated a subset of our framework using Stage [6]. We constructed two 60 m × 60 m USAR settings, each containing significant debris, human victims, and several objects appearing to be human but requiring more capable robots to verify. The robots' goal was to maximize area coverage and correctly identify as many true victims as possible within 30 min. Like previous work [8] we abstracted heterogeneity to three types of robots: a MinBot (intended for exploration and potential victim discovery, but unsuitable as leaders); a MidBot, with advanced victim sensors to correctly identify victims, that can also potentially serve as leaders, and; a MaxBot, having only basic victim sensors but larger memory and computational abilities making them ideal for leadership roles. Teams were composed of 1 MaxBot, 2 MidBots, and 4 MinBots.

We evaluated all three recruitment strategies against three communication success rates (100 %, 60 %, and 20 %), and three probability levels (moderate, low, or no chance) that any robot could experience a random temporary (between 3–4 min.) or permanent failure. To offset these losses, replacement robots (10 MinBots, 2 MidBots, and 1 MaxBot) were released into the environment at the 10-min mark of every trial. Each trial was repeated 50 times in each of our two environments (2700 runs total).

5 Results and Discussion

5.1 Environment Coverage

Results of area coverage (Fig. 1) indicate that poor communication rates encourage greater environmental coverage in active recruitment settings. Failed task assignment attempts lead robots to undertake active searches to find suitable robots to complete tasks which could not be assigned by a leader, resulting in greater coverage. Poor communication further results in failures to recruit other robots, resulting in more searches. Since poor communication results in

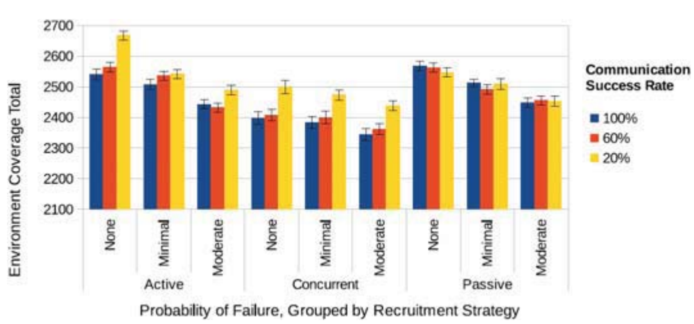

Fig. 1. Sum of all robots' environment coverage, averaged over 50 trials in each configuration. Error bars show standard error.

increased recruitment, exploration tasks are more likely to be assigned to members of different teams as well. This effect is also visible in concurrent settings when communication success rates are low: as recruitment tasks are created to compensate for the lack of assignment acceptances, members of other teams can respond to these requests, resulting in greater coverage of the environment.

5.2 Victims Found

We evaluate our framework against the number of true victims known by a team leader (Fig. 2), since in reality, robots must eventually communicate victim locations to human rescue teams. Team leaders are the logical choice since they are likely to have the most knowledge.

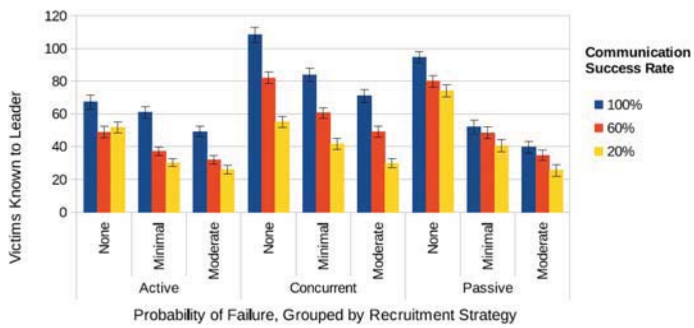

Fig. 2. Sum of all true victims identified and communicated to a team leader, averaged over 50 trials in each configuration. Error bars show standard error.

Concurrent recruitment performs better than other configurations when communication is reliable and robot failures do not occur, since robots continually attempt to recruit others whenever required. This results in a greater number of task assignments overall. In active settings, physical searches for a task are only initiated when a recruitable task is pulled from the robot's task queue. Depending on the priority of the task, there may be a significant delay before recruitment begins. Additionally, active searches take a robot away from its team and leader, resulting in less communication and more lost potential work.

Active recruitment performed poorly in terms of communicating victim locations to a leader. The number of verified victims recorded by individual robots (as opposed to only a leader) indicated that active recruitment results in a high number of victims located when compared to other recruitment strategies, but the results in Fig. 2 indicate that leaders are not well-informed about these victims, since active searches lead a robot away from its leader.

6 Summary and Future Work

We have described a specific subset of recruitment strategies implemented as part of a framework for managing teams of robots in dangerous domains, and evaluated our framework in a simulated USAR environment. Future work will include examining priorities of certain tasks to determine what kind of recruitment strategy, if any, should be used to complete them, as well as the possibility of maintaining copies of useful information over multiple robots to help mitigate the effects of robot failures or insufficient communication with a leader. Information-passing itself could be implemented as a task that would help to improve the effectiveness of the approaches outlined in this paper.

References

1. Brooks, R.A.: A robust layered control system for a mobile robot. IEEE J. Robot. Autom. **2**(1), 14–23 (1986)
2. Costa, E.D., Shiroma, P.M., Campos, M.F.: Cooperative robotic exploration and transport of unknown objects. In: SBR-LARS 2012, pp. 56–61. IEEE (2012)
3. Dos Santos, F., Bazzan, A.L.: Towards efficient multiagent task allocation in the robocup rescue: a biologically-inspired approach. Auton. Agent. Multi-Agent Syst. **22**(3), 465–486 (2011)
4. Dutta, P.S., Sen, S.: Forming stable partnerships. Cogn. Syst. Res. **4**(3), 211–221 (2003)
5. Gage, A., Murphy, R.R.: Affective recruitment of distributed heterogeneous agents. In: AAAI, pp. 14–19 (2004)
6. Gerkey, B., Vaughan, R.T., Howard, A.: The player/stage project: tools for multi-robot and distributed sensor systems. In: Proceedings of the 11th International Conference on Advanced Robotics, vol. 1, pp. 317–323 (2003)
7. Gunn, T., Anderson, J.: Effective task allocation for evolving multi-robot teams in dangerous environments. In: Proceedings of the IAT 2013. IEEE Computer Society Press (2013)
8. Gunn, T., Anderson, J.: Dynamic heterogeneous team formation for robotic urban search and rescue. J. Comput. Syst. Sci. **81**(3), 553–567 (2015)
9. Kiener, J., Von Stryk, O.: Cooperation of heterogeneous, autonomous robots: a case study of humanoid and wheeled robots. In: Proceedings of the IROS 2007, pp. 959–964. IEEE (2007)
10. Krieger, M.J., Billeter, J.B.: The call of duty: self-organised task allocation in a population of up to twelve mobile robots. Robot. Auton. Syst. **30**(1), 65–84 (2000)
11. Mathews, N., Christensen, A.L., O'Grady, R., Rétornaz, P., Bonani, M., Mondada, F., Dorigo, M.: Enhanced directional self-assembly based on active recruitment and guidance. In: Proceedings of the IROS 2011, pp. 4762–4769. IEEE (2011)
12. Pinciroli, C., O'Grady, R., Christensen, A.L., Dorigo, M.: Self-organised recruitment in a heteregeneous swarm. In: ICAR 2009, pp. 1–8. IEEE (2009)
13. Pitonakova, L., Crowder, R., Bullock, S.: Understanding the Role of Recruitment in Collective Robot Foraging. MIT Press, Cambridge (2014)
14. Van De Vijsel, M., Anderson, J.: Coalition formation in multi-agent systems under real-world conditions. In: Proceedings of the AAAI (2004)