Journal of Computer and System Sciences 81 (2015) 553–567

Contents lists available at ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss

Dynamic heterogeneous team formation for robotic urban search and rescue

Tyler Gunn, John Anderson *

Department of Computer Science, E2-445 EITC, University of Manitoba, Winnipeg, MB R3T 2N2, Canada

ARTICLE INFO

Article history:
Received 15 November 2013
Received in revised form 30 May 2014
Accepted 30 July 2014
Available online 18 November 2014

Keywords:
Multi-robot systems
Team formation
Team management
Heterogeneity
Roles
USAR

ABSTRACT

Much work on coalition formation and maintenance exists from the standpoint of abstract agents. This has not yet translated well to robot teams, however: most multi-robot research has focused on pre-formed teams, with little attention to team formation and maintenance. This means solutions fail in challenging environments where equipment is easily lost, such as urban search and rescue. This paper describes a framework for coordinating a changing collection of heterogeneous robots in complex and dynamic environments such as disaster zones. The framework allows a team to reshape to compensate for lost or failed robots, including adding newly-encountered robots or additions from other teams, and also allows new teams to be formed dynamically. The framework also includes provisions for task discovery and assignment, under the conditions of changing team membership. We evaluate this framework through an implementation where robots perform exploration in order to locate victims in a simulated disaster environment.

1. Introduction

Although much research has been performed on teams and coalition formation in multi-agent systems, most works tend to focus on abstract agents performing high level tasks in domains lacking a physical grounding (e.g. package delivery in abstract space [1,2]). While important in principle, these do not take into account many of the physical challenges of being grounded in the real world (e.g. communication distance and reliability [3]). Further, work involving teams of heterogeneous robots is usually restricted to relatively controlled laboratory environments and relies on fixed team structures determined in advance (e.g. [4,5]).

Robots operating in any real-world environment have many challenges to contend with, such as noisy and inaccurate sensor data. Localization is imperfect, and algorithms to intelligently interpret visual data are computationally expensive and inaccurate. Operation in hazardous environments, such as those presented by the exploration of other planets and disaster zones must additionally deal with the fact that robots can be damaged or destroyed. In domains such as these, communication between robots is short range, unreliable, and sporadic in nature: in disaster areas, for example, infrastructure can be heavily damaged, and the debris itself can interfere with wireless communication.

A good example of a highly challenging domain in which robots can be of value is in the aftermath of a natural or man-made disaster. This is commonly known as *Urban Search and Rescue* (USAR), and involves exploring damaged structures to locate and assist human casualties. Operation in a USAR environment presents significant mobility and sensory

---

* Corresponding author.
E-mail address: andersj@cs.umanitoba.ca (J. Anderson).

http://dx.doi.org/10.1016/j.jcss.2014.11.009

difficulties [6]. Debris and uneven terrain can make navigation difficult and can cause a robot to become stuck. Structural changes to the environment as a result of the disaster can render existing floor-plans and maps useless.

The challenges present in a USAR environment make it likely that robots may become lost or separated from their team. Further, robots can become physically damaged or destroyed, impairing the team's effectiveness. It is also possible for different teams of robots operating in geographically separated areas to encounter one another as the mission progresses, providing an opportunity for teams to exchange members or combine resources. New robots can also be expected to arrive sporadically since not all equipment arrives at once or is sent in at the same time.

This paper describes a framework for coordinating a changing collection of heterogeneous robots operating in complex and dynamic environments such as USAR. The framework includes facilities for team formation and management (Section 4.2) as well as facilities for task discovery and assignment (Section 4.3).

2. Related work

Until recently, there has not been a large focus on how to form and maintain teams of robotic agents. Most previous works assume teams were formed in advance and will not change during the course of operation (e.g. [7–12]).

Extensive prior work (e.g. [1,2,13]) has resulted in techniques to enable self-interested agents to form mutually beneficial partnerships in groups of two or more. Although these concepts are generally applicable to robotic domains, these works are demonstrated in domains which are too abstract to show direct applicability for robots in challenging conditions. There is no consideration for issues surrounding the perception of agents in the environment, localization, or the impact of limited range unreliable communication.

George et al. [14] studied a method to form sub-teams in a larger overall team of unmanned aerial vehicles (UAVs) operating in a region. Their approach is more realistic, as it assumes robots are heterogeneous and must cooperate to achieve a common goal. Communication is assumed to be limited range, and operation takes place in a more realistic domain, but there is little or no consideration of the ability to form new teams as opposed to sub-groups.

Cheng and Dasgupta [15] developed a technique to form teams among robots exploring an area. Although their work assumes a more real-world domain, it aims to form teams for the explicit purpose of maximizing the overall explored area, where our work attempts to maintain teams for carrying out a broader set of overall tasks, the nature of which can change over the course of the mission.

Kiener and von Stryk [4] present a framework for the cooperative completion of tasks by teams of heterogeneous robots. Their framework achieves this by modeling the individual tasks of the overall mission, and storing the degree to which each of the robots can perform these tasks. The capabilities of (only) a single humanoid and single wheeled robot are determined in advance, along with weights identifying the suitability of each to all possible tasks. This information allows a central controller to allocate tasks to each robot. While the tasks involved are significant in that they involve fine motion control and interaction, this is still very primitive in terms of task allocation. The broadly different robot skills and task demands result in a predefined set of tasks with only one logical way to map these tasks to the robots in their system. Our framework instead assumes that there may be potentially a large number of potential mappings. Their approach also requires constant communication with a central controller, where our framework performs task allocation in a distributed manner.

Howard et al. [5] developed a system to automatically deploy a sensor network using heterogeneous robots. Resourceful *leader* robots guide network deployment and provide guidance to sensor nodes to keep them in formation. In contrast to our work, the teams, formation, and deployment positions of the sensor nodes are all pre-computed in advance and rely on reliable communication to a central processing unit. Changes in team structure due to loss of robots are accounted for by looking up a new pre-computed deployment pattern and adjusting the formation. No attempt is made to recover lost robots. Task allocation also relies on fixed mapping betweens tasks and robots.

Dorigo et al. [16] developed the *swarmanoids* architecture as a means of encouraging research into swarm robotics in real-world domains. In their work, three heterogeneous robot types cooperate to complete the mission of location a book on a shelf and retrieving it. The capabilities of the robot types, however, preclude them from being used in any other combination, and leaves little opportunity to adapt to changes in available robot types.

Auction-based approaches (e.g. [17,18]) have been developed to perform distributed task allocation amongst teams of heterogeneous robots. Similar to our work, these approaches assume robots can fail at any time and that communication may be unreliable. Auction-based approaches assume bidders will bid only on tasks they are capable of carrying out, where our work makes use of roles to guide the task assignment process to robots which are potentially capable of carrying out a task. Further, auction-based approaches typically assume all robots have the necessary capabilities to assign tasks, where our work assumes task allocation is itself a task which is delegated to a more capable robot.

Gage et al. [19] developed an approach to multi-robot task allocation that uses an emotion-based approach to assigning tasks. Similar to our approach, their work assumes unreliable communication and that robots can fail at any time. In their work, robots continually announce tasks requiring assignment. Those robots that hear the tasks calculate a *shame* value corresponding to their suitability to carry out the task. Similar to a threshold in a market-based system, the shame value determines whether the robot responds with an offer to carry out the task. In contrast to a standard market-based approach, not responding increases the shame value, changing the response threshold. This results in the best suited robots responding first, and the less-suited robots responding later. Although their approach attempts to reduce communication overhead, it places the burden of task allocation on all robots, even the most primitive. This is unrealistic for dangerous environments,