

Exploiting Heterogeneity in Robot Teams Through a Formalism of Capabilities

Wil Thomason and Ross A. Knepper
Department of Computer Science, Cornell University
{wbthomason, rak}@cs.cornell.edu

I. INTRODUCTION

Robot teams are increasingly prevalent, offering potential benefits over solitary robots in scalability and robustness. Solving problems with teams of heterogeneous robots also enables simpler robot design, as team members may be designed for a single task rather than being generally capable. Taking advantage of these benefits requires both *exploiting* and *abstracting* heterogeneity in team members. Exploiting heterogeneity means using the heterogeneous capabilities of team members optimally (i.e. maximizing task completion efficiency and effectiveness); abstracting heterogeneity means hiding all but the minimal necessary set of information about a team of robots. We propose a *formalism of capabilities* that provides heterogeneity exploitation and abstraction for the control and coordination of heterogeneous robot teams.

II. CAPABILITY FORMALISM

Robots in heterogeneous teams need to communicate, make plans, and coordinate in executing these plans. Each of these actions should happen with minimal human involvement. To this end, we propose the following formalism.

We model a capability as a set of effectful function signatures, parameterized on other capabilities or properties of the environment. Any robot which “has” a given capability must provide implementations for each of these functions. This formalism is very similar to the notion of a *typeclass* from the field of programming languages (PL).

This formalism enables heterogeneity exploitation (Section III) and abstraction (Section IV). However, it has other noteworthy properties. First, it lets us borrow many conceptual tools from type theory in PL to reason automatically about capabilities. In particular, we can *automatically compose* capabilities to form novel capabilities that may involve several interacting agents, expressing complex team behaviors. Second, we can prove guarantees about capability use.

III. EXPLOITING HETEROGENEITY

A system for coordinating heterogeneous robot teams must be able to exploit the heterogeneity of those teams. Leveraging information about the team’s heterogeneity produces better results than heterogeneity-agnostic coordination (i.e. guaranteeing task completion whether or not heterogeneity is present). Information about team member capabilities can add useful constraints to the task assignment problem, and a system which incorporates more information always has the option of ignoring it and using the naive strategy.

Our proposed formalism represents a team by its capabilities. This representation explicitly frames a task in terms of uses of capabilities, thus allowing a coordination system to incorporate knowledge about which robots are best suited to which parts of a task in performing task assignment.

IV. ABSTRACTING HETEROGENEITY

Heterogeneity in a robot team can have downsides. Robots can use different control schemas, complete tasks differently, and differ in other ways that make coordination difficult. This difficulty is one of the core limitations of the use of heterogeneous robot teams, as it complicates the programming of teams and restricts the reusability of coordination frameworks.

In order to coordinate a heterogeneous team, a system must have access to information about what actions the team members can take. It also needs access to some notion of control for the team members to make them act. Our formalism of capabilities provides this information — capabilities are modeled as functions, which gives detail about what actions members can take and makes commanding team members as simple as function application. However, our formalism hides information irrelevant to a human coordinating a heterogeneous team by ignoring (in the formalism) the specific mechanisms by which robots implement a capability.

This careful elision of information allows our formalism to be used to create coordination systems which can scale to arbitrary teams of arbitrarily many robots. All that is required is that the team has the capabilities necessary to complete a task. The human programmer can even be unaware of the makeup of the team while writing their code. Thus, by abstracting certain aspects of the heterogeneity of a team, our system enables more broadly useful and adaptable coordination systems.

V. CONCLUSION

Our proposed formalism of capabilities provides the theoretical underpinnings for a system for coordinating heterogeneous robot teams that is both effective (via exploitation of heterogeneity) and flexible (via abstraction of heterogeneity). We are working on a coordination system (called `borg`) capable of adapting to ad hoc heterogeneous teams at arbitrary scale and able to automatically synthesize new behaviors with verifiable guarantees from existing capabilities. We hope that `borg` will be of both practical and theoretical interest to the multi-agent systems community.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. 1646417.