Abstract

With increasing deployment of multiple agents in complex, dynamic settings, there is an increasing need to respond to failures that occur in the agents' coordination. In particular, there is need to detect and diagnose coordination failures—failures to maintain relationships between agents. We refer to this type of diagnosis as social diagnosis. Previous approaches to diagnosis in multi-agent settings have either ignored failures in coordination, or utilized heuristic approaches, which do not scale up as the number of agents (and their interactions) increases.

This dissertation offers a comprehensive and principled approach to social diagnosis. We use a model-based diagnosis (MBD) approach. Here, a model of a diagnosed system is used to simulate the behavior of the system given the operational context (typically, the system's inputs), and to pinpoint possible failing components within the system. MBD has been difficult to apply to diagnosing coordination failures, because of the challenges in constructing a model of coordination, and the lack of appropriate (scalable) diagnosis algorithms.

In the first part of this dissertation, we formally show how to construct a model of agent coordination, and use it to formally define the two key variant social diagnosis problems: the Consistency-based diagnosis problem, and the abductive diagnosis problem. We show that these are NP-Hard problems. We then build on known methods in constraint-satisfaction problems, to provide several algorithms for social diagnosis in centralized and distributed settings. The algorithms—whose analytical guarantees vary in terms of completeness and correctness—are evaluated empirically, in experiments carried out with teams of physical robots. We examined the computational requirements of the algorithms (i.e., their run-time and bandwidth usage), and the correctness of the diagnoses produced. We find that in general, a trade-off exists between computational costs and the correctness of the diagnosis.

In the second part of the dissertation, we focus on a particular type of coordination failures—disagreements—that are of particular interest when talking about teams. We examine the design space of disagreement diagnosis algorithms for a more complex class of situated agents (compared to the first part). We distinguish two phases of diagnosis: (i) the selection of the diagnosing agents; and (ii) the diagnosis of the global team state (by the selected agents). We provide alternative algorithms for these phases, and empirically evaluate their communications and run-time requirements. The results show that centralizing the disambiguation process is a key factor in improving communications, but is not a determining factor in run-time. On the other hand, explicit reasoning about the different agents is a key factor in determining run-time.

Based on this conclusion, we follow two principles in reducing the communications and the computations in large-scale teams. First, we modify the algorithms such that they communicate partial information that is most relevant to the diagnosis. Second, we enable diagnosis of only a limited number of representative agents instead of diagnosing all others. These principles yield a novel diagnosis method that significantly reduces the runtime, while keeping communications overhead to a minimum.

A Study of Dynamic Coordination Mechanisms

Candidate: Avi Rosenfeld
Institution: Department of Computer Science, Bar Ilan University, Israel
Supervisors: Gal A. Kaminka and Sarit Kraus

Year awarded: 2007
URL: http://www.jct.ac.il/~rosenfa/
doi:10.1017/S0269888909990294

Abstract

Coordination, or the act of managing interdependencies between activities, is a key issue within the field of multi-agent systems. Because of the importance of this issue, many theoretical and practical frameworks have been proposed for addressing coordination challenges. However, finding the optimal coordination method for a given group of agents a domain task is a computationally difficult, if not intractable, problem in most real-world domains. Solving the 'coordination problem' is thus an important open challenge for researchers in this field.

In a variety of domains, various coordination algorithms have been developed to date, with each algorithm being most successful for different domain conditions. First, we address how one can create adaptive coordination methods from existing coordination methods for robots without communication. We then extend these results to address adaptation between coordination algorithms for robotic groups using communication. Third, we investigate methods for creating dynamic coordination methods from distributed constraint optimization algorithms. Finally, we study various coordination algorithms within distributed peer-to-peer search algorithms.

Across these diverse domains, this thesis presents a unified algorithm selection approach for creating effective adaptive coordination methods. We find that novel teamwork measures can be developed for quantifying the effectiveness of coordination algorithms in all of these domains. These measures can be autonomously and locally measured by team members, even without any communication. The significance of this result is its ability to effectively quantify coordination in a clear, tractable fashion. Next, we find that these measures can be used to switch between coordination methods as needed. Robots or agents can effectively select the best coordination method to their localized domain conditions, online during task execution. The net result is a significant productivity improvement of these adaptive methods over the static methods they are based on.