Diagnosing Multi-agent Systems

Meir Kalech and Gal A. Kaminka

The MAVERICK Group
Computer Science Department
Bar-Ilan University, Israel
{kalechm,galk}@cs.biu.ac.il

One of the key requirements in collaborating distributed applications is that their sub-systems remain synchronized during their joint operation. With increasing deployment of software "teams", composed of distributed applications in complex, dynamic settings, there is an increasing need to also be able respond to failures that occur in their synchronization and coordinated operation, in particular to be able to diagnose the causes for disagreements (synchronization failures) that may occur, in order to facilitate recovery and reestablish collaboration e.g., by automated negotiations. This type of diagnosis is called social diagnosis, since it focuses on finding causes for failures to maintain designer-specified relationships between sub-systems. Naive implementations of social diagnosis processes can require significant computation and communications overhead, which prohibits them from being effective as the number of sub-systems is scaled up, or the number of failures to diagnose increases. We thus seek to examine in depth the communication and computation overhead of diagnosis.

For instance, suppose a company has multiple sub-systems that must interact with each other to achieve a common task, e.g., two cellular phone stations (cells) that must pass a cellular phone call from one to another as the user moves. The two stations must coordinate on the frequencies involved and other salient information. Sometimes, due to failures in communications between the two stations, or due to intermittent faults in the system, this coordination fails, and the two stations remain un-coordinated. Discovering the problem can fairly easy (the call gets disconnected), but determining the cause for it (failed communication links between the stations) can be very difficult. One reason for this is that the fault may lie at a cause that is in-between the stations, such as the communication medium.

We treat each sub-system (in the example, each station) as an agent, and the system as a whole (with its many distributed sub-systems as an organization, or a team). We distinguish two phases of diagnosis: (i) selection of the diagnosing agents; and (ii) diagnosis of the global team state (by the selected agents). We provide alternative algorithms for these phases, and combine them in different ways, to present diagnosis methods, corresponding to different design decisions.

We empirically evaluated the communications and run-time requirements of these methods in diagnosing thousands of systematically-generated failure cases, occurring in a complex simulation application. The results show that centralizing the disambiguation process is a key factor in dramatically improving communications efficiency, but is not a determining factor in run-time efficiency. On the other hand, explicit reasoning about the different sub-systems is a key factor in determining run-time: Methods that require explicit reasoning about different sub-systems incur significant computational costs, though they are sometimes able to reduce the amount of communications.

Based on the conclusion that centralizing the diagnosis reduces the communication, we address two principles to achieve the reduction of the communication and the computation in teams where the number of agents is scaled-up. First, we disambiguate the information sent by the agent before communicating: Instead of sending all the information, send only the information that is relevant to the diagnosis. Second, we diagnose a limited number of agents that represent all others. These principles yield a novel diagnosis method which identifies the minimal number of agents that are necessary in order to make the diagnosis process. This method significantly reduces the runtime, while keeping communications overhead to a minimum.

References

[Kalech and Kaminka2003] Meir Kalech and Gal A. Kaminka. On the design of social diagnosis algorithms for multi-agent teams. *IJCAI-03*, pages 370–375, 2003.

[Kaminka and Tambe2000] Gal A. Kaminka and Milind Tambe. Robust multi-agent teams via socially-attentive monitoring. *Jornal of Artificial Intelligence Research*, 12:105–147, 2000.