

Task Coordination in Multiagent Systems

Extended Abstract

Vahid Yazdanpanah
University of Twente
Enschede, The Netherlands
V.Yazdanpanah@utwente.nl

Mehdi Dastani
Utrecht University
Utrecht, The Netherlands
M.M.Dastani@uu.nl

Shaheen Fatima
Loughborough University
Loughborough, United Kingdom
S.S.Fatima@lboro.ac.uk

Nicholas R. Jennings
Imperial College London
London, United Kingdom
N.Jennings@imperial.ac.uk

Devrim M. Yazan
University of Twente
Enschede, The Netherlands
D.M.Yazan@utwente.nl

Henk Zijm
University of Twente
Enschede, The Netherlands
W.H.M.Zijm@utwente.nl

ABSTRACT

We present an overview of the Task Coordination (TC) problem in multiagent systems and discuss the specific elements that are required to develop a solution to this problem. Task coordination refers to a twofold problem where an exogenously imposed state of affairs should be satisfied by a multiagent system (MAS): (1) the agents need to be assigned tasks to fulfill the given state of affairs (*task allocation*) and (2) the behavior of agents needs to be monitored to evaluate whether their tasks are fulfilled so that responsibility for dismissed tasks can be determined (*task responsibility*). This becomes especially challenging when agents are autonomous and may have imperfect information about their environment. Then, the allocation of tasks and responsibilities should regard agents' strategic ability under imperfect information. To date, existing work on the application of strategic reasoning for task allocation assumes perfect information for agents (dismissing imperfect information settings) and allocates tasks to individual agents (dismissing the potential for allocating tasks to agent groups). This calls for TC frameworks able to model task allocation in imperfect information settings and by allowing the allocation of tasks to agent groups. Such a framework should also be able to determine the responsibility of agents for dismissed tasks via a task responsibility mechanism that complements the task allocation procedure. This work discusses various aspects of the TC problem, sets forward a conceptual analysis on expected properties of potential solution concepts, and presents the overview of a suggested approach for developing a TC framework using techniques from formal strategic reasoning.

KEYWORDS

Task Coordination in MAS, Task Allocation, Responsibility Reasoning, Strategic Reasoning, Applied Formal Methods

ACM Reference Format:

Vahid Yazdanpanah, Mehdi Dastani, Shaheen Fatima, Nicholas R. Jennings, Devrim M. Yazan, and Henk Zijm. 2020. Task Coordination in Multiagent Systems. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020)*, Auckland, New Zealand, May 9–13, 2020, IFAAMAS, 3 pages.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1 INTRODUCTION

Task Coordination (TC) is a twofold problem in Multiagent Systems (MAS). Given a state of affairs—as an exogenously imposed set of properties to, and expected to be fulfilled by, a MAS—it is crucial to have a systematic method for *allocating tasks* to agents (prospectively) and *ascribing responsibilities* to agents based on what they were tasked to do and what they actually did (in retrospection). This way, TC consists of two stages: Task Allocation (TA) and Task Responsibility (TR). TA is concerned with how the state of affairs should be distributed among agent groups in terms of *tasks*. Then TR is about evaluating the behavior of the MAS in fulfilling the tasks and ascribing (a degree of) *responsibility* to agents for dismissed tasks. We believe that ascribing responsibility to agents is justified only if the task allocation process takes into account the *strategic abilities* of agents and their *epistemic limitations*. In brief, strategic abilities determine what agents can do (in terms of properties they can ensure/preclude) while epistemic limitations are about their potential (lack of) knowledge about the MAS. In most real-world environments, the task allocation procedure should incorporate these aspects.

Capturing strategic abilities under imperfect information implies that a task should be allocated to agents or agent groups that are capable of handling it. It is crucial to note that agents' ability is limited to their knowledge about the environment. Thus it is necessary to capture strategic abilities in imperfect information settings and to avoid assuming perfect information for all agents. Moreover, there might be tasks for which no *single* agent has the required capabilities. Then it is necessary to allow the allocation of such tasks to capable agent *groups*—instead of dismissing them. A natural approach to capture these aspects is to apply strategic reasoning and formal methods [1]. However, no such method currently exists [5, 11]. In [5], authors capture *strategic* abilities but under a perfect information assumption while [11] relaxes this assumption for *team formation*—as a related problem to TC—but has less temporal expressivity than [5].

TR, as the retrospective aspect of the TC problem, is about the consideration that agents may fail to deliver an allocated task. Thus it is crucial to monitor the history of materialized actions and ascribe responsibility to agent groups—accordingly to individuals—based on what they did versus what they had to do.

Against this background and to capture the real-life subtleties of the TC problem, we see the need for development of a TC framework that: (1) integrates TA and TR, (2) captures strategic abilities under imperfect information, and (3) enables the allocation of tasks and responsibilities to agents as well as agent groups.

2 ESSENTIAL AND DESIRABLE PROPERTIES

We first identify major conceptual principles for task coordination in multiagent systems.

- *Suitability* of the state of affairs: Given the set of agents and their available actions, fulfilling some states of affairs are impossible in principle, regardless of how we allocate tasks among the agents. We deem that for TC, one can simply exclude this class of unsuitable and out of reach expectations. Then we can assume that a state of affairs is *suitable* only if the grand coalition (of agents) is able to ensure it.
- *Validity* of task allocation: Task allocation ought to be such that all that should be done is allocated, neither more nor less. In other words, we see an allocation of tasks as *valid* if by assuming that all agents fulfill their allocated tasks, the state of affairs will ensure.
- *Justifiability* of task responsibility: Task responsibility should be consistent with task allocation. In other words, seeing a group as being responsible is not independent of what tasks they were given in an earlier stage. Thus, ascribing task responsibility is not merely based on agents' ability but has to build upon the implemented task allocation and the history of realized actions.

It is essential that a TC framework fulfils these basic principles. In principle, for a *suitable* state of affairs, a task allocation procedure should ensure a *valid* allocation of tasks. In other words, it should distribute tasks such that the state of affairs (as a whole) can be fulfilled. Then the responsibility ascription component should build upon the task allocation procedure to regard the justifiability principle. In brief, while the task allocation component is about eventualities, the responsibility ascription should consider the actuality of agents' realized acts—in comparison to what they were tasked to do. To ascribe responsibility in multiagent settings, one may face the so-called “*responsibility gaps*” [10] where a group is responsible but the extent of each member's responsibility is unclear. In such a situation, an effective TC framework should be able to determine a *degree* of responsibility for group members.

In addition to the basic principles that a TC framework should satisfy, we see some properties as desirable in specific settings. For instance in domains with a high level of uncertainty, *fault tolerance* is desirable. Then one can allocate a task to a group and also introduce a set of backup groups. See [4] for a suggestion to employ responsibility reasoning for ensuring fault tolerance. Another concern is the *efficiency* of a task allocation. This is to rank task allocations based on their cost and take into account the resources they consume (e.g., time). Finally, it is desirable to have a *strategy-proof* TC framework. This is to ensure the compatibility of task allocations with agents' preferences by guaranteeing that agents have no incentive to deviate from what they are tasked to do. For this, we can rely on normative incentivization mechanisms [3].

3 A SUGGESTED APPROACH

A suggested approach for developing an effective TC framework is to use the semantic machinery of temporal multiagent logics, in particular *Concurrent Epistemic Game Structures (CEGS)* [1]. (CEGS is the epistemic extension of Concurrent Game Structures [2].) In addition to being expressive for specifying temporal, strategic, and epistemic aspects of MAS, models that use CEGS can benefit from standard model checking platforms to verify properties of the modeled MAS (e.g., model-checking tools in [7, 9]).¹

In a CEGS-modeled MAS, a state of affairs can be modeled as a set of formulae that represents some properties that are expected to hold in (some) CEGS states. Then the task allocation component should ensure that all the members of the set are allocated to agents (a forward-looking procedure). However, having autonomous agents, we may see some tasks remaining unfulfilled. Then the task responsibility component ascribes a degree of responsibility to agents for dismissed tasks (a backward-looking procedure).

As discussed earlier, task allocation should allow the allocation of tasks to agent groups (and not merely to individuals) and also relax the assumption that agents have perfect information about their environment and the consequence of their actions. Thus we can allocate the task of ensuring a formula to an agent group only if they are both strategically and epistemically capable of fulfilling it. For this, the CEGS-modeled notion of “*uniform strategy*” [8] is applicable as it allows reasoning about capabilities of agents and groups under imperfect information. Assuming that the allocation of tasks to agents definitely brings about the given state of affairs is unreasonable in real-life environments. (Agents are not artifacts but entities that may opt to exercise their autonomy and do other than what they are tasked to do.) A TC framework should be conscious of this and ascribe a degree of responsibility to agents, in particular to those who contributed to a state of affairs remaining unfulfilled. To that end, one can adopt the notion of “*strategic responsibility*” form [12] as a CEGS-based responsibility reasoning method for imperfect information settings.

4 CONCLUSION

We discussed TC as a multidimensional problem in MAS, argued for properties that are expected from a potential solution concept, and suggested an approach to apply formal strategic reasoning methods. In future work, we aim to explore the suggested CEGS-based approach and develop a TC framework that satisfies the essential properties and enables the integration of further desirable aspects. We are also interested in exploring the dynamics of TC in various organizational structures [6].

Acknowledgments: This work is a result of the first author's research visit at Imperial College London under the support of the University of Twente. He acknowledges the supports provided by both institutes. The authors would also like to thank the anonymous referees for their valuable comments and helpful suggestions.

¹In CEGS, a MAS is modeled as a transition system. In an informal language, in each state of this transition system, some properties hold and agents are able to change the state of the system (hence its properties) by executing their available actions. Moreover, we have a representation of each agent's epistemic limitation in terms of system states that are indistinguishable to her.

REFERENCES

- [1] Thomas Ågotnes, Valentin Goranko, Wojciech Jamroga, and Michael Wooldridge. 2015. Knowledge and Ability. In *Handbook of Epistemic Logic*. College Publications, London, UK, 543–589.
- [2] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. 2002. Alternating-time temporal logic. *J. ACM* 49, 5 (2002), 672–713. <https://doi.org/10.1145/585265.585270>
- [3] Nils Bulling and Mehdi Dastani. 2016. Norm-based mechanism design. *Artif. Intell.* 239 (2016), 97–142. <https://doi.org/10.1016/j.artint.2016.07.001>
- [4] Hana Chockler and Joseph Y. Halpern. 2004. Responsibility and Blame: A Structural-Model Approach. *J. Artif. Intell. Res.* 22 (2004), 93–115. <https://doi.org/10.1613/jair.1391>
- [5] S. Shaheen Fatima and Michael Wooldridge. 2001. Adaptive Task Resources Allocation in Multi-Agent Systems. In *Proceedings of the Fifth International Conference on Autonomous Agents (AGENTS '01)*. Association for Computing Machinery, New York, NY, USA, 537–544. <https://doi.org/10.1145/375735.376439>
- [6] Bryan Horling and Victor R. Lesser. 2004. A survey of multi-agent organizational paradigms. *Knowledge Eng. Review* 19, 4 (2004), 281–316. <https://doi.org/10.1017/S0269888905000317>
- [7] Wojciech Jamroga and Jürgen Dix. 2005. Model Checking Strategic Abilities of Agents Under Incomplete Information. In *Theoretical Computer Science, 9th Italian Conference, ICTCS 2005, Siena, Italy, October 12-14, 2005, Proceedings (Lecture Notes in Computer Science)*, Vol. 3701. Springer, Berlin, Germany, 295–308. https://doi.org/10.1007/11560586_24
- [8] Wojciech Jamroga and Wiebe van der Hoek. 2004. Agents that Know How to Play. *Fundam. Inform.* 63, 2-3 (2004), 185–219. <http://content.iospress.com/articles/fundamenta-informaticae/fi63-2-3-05>
- [9] Alessio Lomuscio, Hongyang Qu, and Franco Raimondi. 2017. MCMAS: an open-source model checker for the verification of multi-agent systems. *STTT* 19, 1 (2017), 9–30. <https://doi.org/10.1007/s10009-015-0378-x>
- [10] Andreas Matthias. 2004. The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and information technology* 6, 3 (2004), 175–183.
- [11] Ranjit Nair, Milind Tambe, and Stacy Marsella. 2002. Team Formation for Reformation in Multiagent Domains Like RoboCupRescue. In *RoboCup 2002: Robot Soccer World Cup VI (Lecture Notes in Computer Science)*, Vol. 2752. Springer, Berlin, Germany, 150–161. https://doi.org/10.1007/978-3-540-45135-8_12
- [12] Vahid Yazdanpanah, Mehdi Dastani, Wojciech Jamroga, Natasha Alechina, and Brian Logan. 2019. Strategic Responsibility Under Imperfect Information. In *Proceedings of the 18th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS '19)*. International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, 592–600. <https://dl.acm.org/doi/10.5555/3306127.3331745>