

Multiagent Task Coordination as Task Allocation Plus Task Responsibility

Vahid Yazdanpanah¹, , Mehdi Dastani², Shaheen Fatima³,
Nicholas R. Jennings⁴, Devrim M. Yazan⁵, and Henk Zijm⁵

¹ University of Southampton, Southampton, United Kingdom

V.Yazdanpanah@soton.ac.uk

² Utrecht University, Utrecht, The Netherlands

M.M.Dastani@uu.nl

³ Loughborough University, Loughborough, United Kingdom

S.S.Fatima@lboro.ac.uk

⁴ Imperial College London, London, United Kingdom

N.Jennings@imperial.ac.uk

⁵ University of Twente, Enschede, The Netherlands

{D.M.Yazan, W.H.M.Zijm}@utwente.nl

Abstract. In this work, we present a dynamic Task Coordination framework (TasCore) for multiagent systems. Here task coordination refers to a twofold problem where an exogenously imposed state of affairs should be satisfied by a multiagent system (MAS). To address this problem the involved agents or agent groups need to be assigned tasks to fulfill (task allocation) and the behavior of these agents needs to be monitored to evaluate whether their tasks are fulfilled so that responsibility for dismissing tasks can be determined (task responsibility). We believe the allocation of tasks should regard both the strategic abilities of agents and their epistemic limitations. To date, however, existing work on the application of logical strategic reasoning for task allocation assumes perfect information for agents (dismissing imperfect information settings) and allocates tasks to individual agents (dismissing task allocation to agent groups). In TasCore, we address this gap by modeling *task allocation* using imperfect information semantics for strategic reasoning and integrate it with a notion of *task responsibility*. We formally verify properties of TasCore: on validity as well as stability of task allocations and fairness as well as non-monotonicity of task responsibilities.

Keywords: Task Coordination in Multiagent Systems · Responsibility and Accountability · Strategic Reasoning · Temporal and Modal Logic.

1 Introduction

The focus of this paper is on the *task coordination (TC)* problem in Multiagent Systems (MAS). Given a state of affairs (exogenously imposed to, and to be fulfilled by, a MAS) it is crucial to have a systematic method for *allocating tasks* to involved 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). Given a collective state of affairs (to be fulfilled by the MAS) 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 dismissing/fulfilling tasks. In other words, following the allocation of a task to an agent or agent group, we see the group as being responsible for fulfilling the allocated task. In a MAS, there might be tasks for which no *single* agent has the required capabilities. Our work allows such tasks to be allocated to capable agent *groups* (instead of dismissing the task).

We believe that ascribing responsibility to agents is justified only if the task allocation process takes into account the *strategic abilities* of the agents and their *epistemic limitations*. In brief, strategic abilities of agents or agent groups determine what they can do in the MAS (e.g., in terms of properties that they can ensure or preclude), while epistemic limitations are about their potential (lack of) knowledge about the MAS. Capturing these two aspects for allocating tasks and ascribing responsibilities in MAS are focal points of this work.

In general, the TC problem relates to studies on Multiagent Organizational (MAO) frameworks, task allocation methods in MAS, and responsibility reasoning concepts in multiagent settings. Reviewing MAO frameworks in [13, 32, 21, 24], their focus is on how the MAO is organized in terms of its organizational structure and high level constructs to enable role/task allocation. They abstract from the exact procedure of task allocation and in some cases assume the availability of agents that are capable of fulfilling any assigned task. In MAO—as a whole or within any of its organizational units—task allocation techniques are often employed as a module to determine who should do what task(s) to ensure the organizational goals (as a collectively defined state of affairs). For instance, given the goal to have a picnic in the countryside for a two-member organization, one agent can be responsible for driving while the other one is responsible for food (i.e., allocating the task to drive to one agent and to prepare the food

to the other one). Then some relevant questions are: “which one is able to drive (*strategic ability*)?”; “who knows about potential food allergies (*epistemic limitations in imperfect information settings*¹)?”; “who is able to cook/drive at what time, as the journey/picnic evolves (*temporal dynamics*)?”. In most real-world environments, the task allocation procedure should capture all three aspects, i.e., addressing strategic, epistemic, and temporal aspects. However, no such method currently exists [16, 9, 3, 11, 26, 29]. In [16], authors capture *strategic* abilities but under a perfect information assumption. The presented approach in [29] relaxes this assumption for *team formation*—as a related problem to TC—but has less temporal expressivity than [16].

Recall that in the picnic example, we said that an agent is *responsible* for food or for driving. That is, when we arrive at the picnic spot with no food, we can justifiably see the agent (to whom we allocated the task of providing food) as being responsible for the (undesirable) outcome. We argue that dealing with autonomous agents, allocating tasks (prospectively) needs to be complemented by a (retrospective) phase of responsibility allocation. In TAsCore, the TR component answers this question by ascribing a degree of responsibility to each agent, for an outcome φ . This means seeing an agent as being responsible (for φ) to a degree proportional to its contribution to the occurrence of φ . To model task responsibility, we apply multiagent responsibility reasoning methods [2, 35], in particular the notion of *strategic responsibility* [36], for ascription of responsibility in such imperfect information settings.

To overcome these shortcomings, we develop TAsCore: a dynamic task coordination method based on formal methods for multiagent strategic reasoning. TAsCore is a two-part method: its prospective part is focused on task allocation, while the retrospective part is focused on task responsibility. For task allocation, we allocate a task to an agent or agent group that is capable of handling it. One aspect that we see as being crucial to capture is the fact that the agents’ ability is limited to their knowledge about the environment. It is therefore necessary to capture strategic abilities in imperfect information settings and to avoid assuming perfect information for all agents. In imperfect information settings, agent A may be able to guarantee the fulfillment of a task τ

¹ By *imperfect information* settings, we refer to the generic class of multiagent settings in which agents do not (necessarily) have perfect information about their environment and all the potential consequences of their actions (also referred to as agents with *epistemic limitations*). In most real-life applications of MAS, agents are epistemically limited. Thus, it is crucial to develop a model that is capable of capturing imperfect information settings.

today, but not necessarily tomorrow—due to the information and strategies that it possesses today and may miss tomorrow². This is a missing aspect in most task allocation methods for MAS. It also motivates our approach to use the semantic machinery of multiagent logics that are expressive for modeling temporal and strategic properties in imperfect information settings. The retrospective part of TASCORE is focused on the history and is about ascribing responsibility to agents—considering what they did and what they had to do. As we allocate tasks by taking into account agents’ abilities, the fulfillment of an allocated task is a justified expectation, hence agents that violate this expectation are justifiably responsible for it. Against this background, for the first time, this paper captures strategic abilities *under imperfect information* for task allocation in multiagent settings and applies *responsibility reasoning* for task coordination in organizational settings.

We present a conceptual analysis on the TC problem and recall formal preliminaries in Section 2. Then, in Section 3, we specify TASCORE and its components. In Sections 4 and 5, we focus on the allocation of tasks and ascription of responsibilities, respectively. We formally verify properties of TASCORE: on validity as well as stability of task allocations and fairness as well as non-monotonicity of task responsibilities. Finally, Section 6 discusses the relevance and implementability of TASCORE and presents the concluding remarks.

2 Conceptual Analysis and Formal Preliminaries

In this section, we present the intuition behind our work, analyze conceptual aspects of TC, and recall key formal notions that form the basis of the TASCORE framework.

2.1 Conceptual Analysis

Imagine a family picnic scenario in which parents allocate different picnic-related tasks to their children *Alex*, *Bob*, *Cathy*, and see them responsible for organizing a picnic that satisfies some expectations. They expect the picnic to be organized in a clean picnic spot in the country side and to have home-made

² For instance, an agent may be perfectly capable of planning a chain of actions to prepare a meal at home—given all the equipment—but not able to do so if she moves to a picnic camp the day after, e.g., due to inconvenient weather and the lack of equipment.

food for the day. In this case, having a desirable picnic is a collective state of affairs. Then the first bit of the task coordination problem is to see “*who is able to do what*” in a temporal, strategic, and imperfect information setting. The second bit is about “*seeing who did what*” and ascribing responsibility to agents for the outcome. We deem that the second phase (task responsibility) is justified only if in the first phase (task allocation), the strategic, temporal, and epistemic aspects of the setting are captured. For instance, if we give the task of driving to the spot to Cathy while she has no driving licence, it is not reasonable to blame her for dismissing the allocated task. The analogous case is valid for giving the task of cleaning to Bob, if he cannot distinguish clean from unclean (e.g., due to a visual impairment). We later show how these aspects can be modeled and verified using the semantic machinery of logics for multiagent strategic reasoning. In addition to these aspects, there are a number of high-level principles that are essential to effective task coordination:

Suitability of the collective 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. For instance, in our picnic scenario, a modest picnic is reasonable to expect but seeing the small group organize a festival might not be a *suitable* (collective) state of affairs.

Validity of task allocation: given a *suitable* state of affairs, the process of task allocation ought to be such that all that should be done is allocated, neither more nor less (i.e., if all agents fulfill their tasks, the collective state of affairs will be fulfilled). For instance, if we tell Alex to take care of driving and give the task of cleaning the place to Bob and Cathy, the allocation is not valid as preparing food is dismissed. Basically, we see an allocation of tasks as *valid* if by assuming that all agents fulfill their allocated task, we can see the collective 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. For instance, if we give the task of preparing food to Alex, then it is not *justifiable* to see Bob responsible for the food quality even if he is able to cook. This shows that verifying task responsibility is not merely based on agents’ ability but builds upon the implemented task allocation and the history of realized actions (i.e., the evolution of the multiagent system).

In subsequent sections, we provide a formal account of these principles and then fulfill them using TasCore . To specify the multiagent setting, we use *Concurrent Epistemic Game Structures* (CEGS) [1] as it allows: modeling the behavior of a MAS, specifying strategic abilities of agents, and representing agents' knowledge. Then we focus on task and responsibility allocation in TasCore .

2.2 Concurrent Epistemic Game Structures

To model multiagent systems and analyze their strategic behavior under imperfect information, we use *Concurrent Epistemic Game Structures* (CEGS) [1] as an epistemic extension of Concurrent Game Structures [4]. 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 (e.g., ATL-based model-checking tools in [25,23]) to verify properties of the modeled MAS.

Concurrent Epistemic Game Structures: Formally, a *concurrent epistemic game structure* is a tuple $\mathcal{M} = \langle \Sigma, Q, Act, \Pi, \pi, \sim_1, \dots, \sim_n, d, o \rangle$ where: $\Sigma = \{a_1, \dots, a_n\}$ is a finite non-empty set of *agents*; Q is a finite non-empty set of *states*; Act is a finite set of atomic *actions*; Π a set of atomic *propositions*; $\pi : \Pi \mapsto 2^Q$ is a propositional evaluation function; $\sim_a \subseteq Q \times Q$ is an *epistemic indistinguishability relation* for each agent $a \in \Sigma$ (we assume that \sim_a is an equivalence relation, where $q \sim_a q'$ indicates that states q and q' are indistinguishable to a); function $d : \Sigma \times Q \mapsto \mathcal{P}(Act)$ specifies the sets of actions available to agents at each state (we require that the same actions be available to an agent in indistinguishable states, i.e., $d(a, q) = d(a, q')$ whenever $q \sim_a q'$); and o is a deterministic transition function that assigns the outcome state $q' = o(q, \alpha_1, \dots, \alpha_n)$ to state q and a tuple of actions $\alpha_i \in d(i, q)$ that can be executed by Σ in q .

To enable the specification of a collective state of affairs, we adopt the standard language of LTL (Linear Temporal Logic [30]). Formulas of the language \mathcal{L}_{LTL} are defined by the following syntax, $\varphi, \psi ::= p \mid \neg\varphi \mid \varphi \wedge \psi \mid \bigcirc\varphi \mid \varphi\mathcal{U}\psi \mid \square\varphi$ where $p \in \Pi$ is a proposition, \neg and \wedge are standard logical operators, $\bigcirc\varphi$ means that φ is true in the next state of \mathcal{M} , $\psi\mathcal{U}\varphi$ means that ψ has to hold at least until φ becomes true; and $\square\varphi$ means that φ is always true. For convenience, $\diamond\varphi$ is defined as an equivalent to $\neg\square\neg\varphi$ meaning that φ is eventually true. To represent and reason about *strategies* and *outcomes* in agent systems with imperfect information, we make use of the following auxiliary notions. (References to ele-

ments of \mathcal{M} are to elements of a CEGS \mathcal{M} modeling a given multiagent system, e.g., we write Q instead of Q in \mathcal{M} .)

Successors and Computations: For two states q and q' , we say q' is a *successor* of q if there exist actions $\alpha_i \in d(i, q)$ for $i \in \{1, \dots, n\}$ in q such that $q' = o(q, \alpha_1, \dots, \alpha_n)$, i.e., agents in Σ can collectively guarantee in q that q' will be the next system state. A *computation* of a CEGS \mathcal{M} is an infinite sequence of states $\lambda = q_0, q_1, \dots$ such that, for all $k > 0$, we have that q_k is a successor of q_{k-1} . We refer to a computation that starts in q as a *q-computation*. For $k \in \{0, 1, \dots\}$, we denote the k 'th state in λ by $\lambda[k]$, and $\lambda[0, k]$ and $\lambda[k, \infty]$ respectively denote the finite prefix q_0, \dots, q_k and infinite suffix q_k, q_{k+1}, \dots of λ . We refer to any two arbitrary states q_k and q_{k+1} as two *consecutive* states in $\lambda[k, \infty]$. Finally, we say a finite sequence of states q_0, \dots, q_n is a *q-history* if $q_n = q$, $n \geq 1$, and for all $0 \leq k < n$ we have that q_{k+1} is a successor of q_k . We denote a *q-history* that starts in q_k and has n steps with $\lambda[q_k, n]$. The set of *q-histories* for all $q \in Q$ is denoted by \mathcal{H} .

Strategies and Outcomes: A *memoryless imperfect information strategy*³ for an agent $a \in \Sigma$ is a function $\zeta_a : Q \mapsto Act$ such that, for all $q \in Q$: (1) $\zeta_a(q) \in d(q, a)$, and (2) $q \sim_a q'$ implies $\zeta_a(q) = \zeta_a(q')$. For a group of agents $\Gamma \subseteq \Sigma$, a *collective strategy* $Z_\Gamma = \{\zeta_a \mid a \in \Gamma\}$ is an indexed set of strategies, one for every $a \in \Gamma$. Then, $out(q, Z_\Gamma)$ is defined as the set of potential *q-computations* that agents in Γ can enforce by following their corresponding strategies in Z_Γ . We extend the notion to sets of states $\chi \subseteq Q$ in the straightforward way: $out(\chi, Z_\Gamma) = \bigcup_{q' \in \chi} out(q', Z_\Gamma)$.

Uniform Strategies: A uniform strategy is one in which agents select the same actions in all states where they have the same information available to them. In particular, if agent $a \in \Sigma$ is uncertain whether the current state is q or q' , then a should select the same action in q and in q' . Formally, a strategy ζ_a for agent $a \in \Sigma$ is called *uniform* if for any pair of states q, q' such that $q \sim_a q'$, $\zeta_a(q) = \zeta_a(q')$. A strategy Z_Γ is uniform if it is uniform for every $a \in \Gamma \subseteq \Sigma$. Realistic modeling of strategic ability under imperfect information requires restricting attention to uniform strategies only.

³ We relax the assumption that agents have a perfect memory. Thus, as a natural choice in imperfect information settings, we focus on memoryless strategies and avoid other forms of strategy that assume the ability of agents to recall the evolution of the MAS, e.g., perfect recall strategies (see [8]). This captures a more generic class of agents as we do not expect them to have the capacity to strategize based on their memory of events.

3 Specification

To specify `TasCore`, four components are required: a behavior modeling machinery, a collective state of affairs (given to the MAS), the task allocating component, and the task responsibility component.

Given a CEGS $\mathcal{M} = \langle \Sigma, Q, Act, \Pi, \pi, \sim_1, \dots, \sim_n, d, o \rangle$ that models the behavior of the multiagent system, a collective state of affairs G_q (given to MAS in state q) is a set of formulae from \mathcal{L}_{LTL} .⁴ Then the aim of the task allocation process is to ensure that all $\varphi \in G_q$ hold. Finally, the task responsibility process ascribes a (backward-looking) degree of responsibility to agents given a history $h \in \mathcal{H}$. In this approach, \mathcal{M} is a standard component (adapted from [1]) for modeling the behavior of a MAS in imperfect information settings. Then G_q specifies the set of properties that are expected to be satisfied by the agents collectively (we call it the collective state of affairs). Then what each subgroup ought to do is determined by *task allocation* and who is responsible to what extent by *task responsibility*. We present the exact specification of both in upcoming sections. This means `TasCore` is built on a behavior-modeling CEGS, a local state of affairs, and allocates (forward-looking) tasks as well as (backward-looking) responsibilities to agent groups. Note that all the elements of `TasCore` are defined independently of any desirable properties. So, how an element should be specified such that desirable properties emerge is not intrinsic to the model but will be discussed in the following sections. This is to allow a level of flexibility and to enable capturing context-dependent concerns. `TasCore` supports task coordination using two forms of prospective and retrospective organizational reasoning. The former is about allocating tasks to agents (Section 4) while the latter is about verifying what went wrong/right and who is responsible to what extent (Section 5).

⁴ One may opt to specify the collective state of affairs simply using propositions from Π . However, this will limit the expressivity. Our LTL formulas enable the specification of temporally-bounded tasks. E.g., simply giving a task to ensure that food is ready may lead to giving the task to someone with a strategy to ensure it in the next week while we aim to have the picnic during the weekend (e.g., this can be achieved by ensuring that food is prepared before the weekend using the “*until*” modality in LTL). Moreover, for task specification, it is necessary to be able to express a dynamic behavioral (semantic) form of task (which is more expressive than a static propositional (syntactic) notion of a task). For instance, one may give (an agent) the task of not only ensuring the tidiness of a place but also to maintain it. Such a notion of task is not about ensuring a state or a set of states (where tidiness holds) but about ensuring a chain/path/computation of states (through which the tidiness is maintained). In general, temporal modalities of LTL enable capturing such temporal subtleties.

4 Allocating Tasks in TasCore

Following the specification of TasCore components, and given a local state of affairs G_q (to be fulfilled collectively by agents in the MAS), “*who should do what, and why?*” is the main question that we aim to answer in this section⁵. We deem that the ascription of tasks to agents or agent groups (with the aim to fulfill a collective state of affairs) should take into account the *temporal*, *strategic*, and *epistemic* aspects of multiagent systems. Temporality is both about the specification of tasks (e.g., whether some property should be maintained or only achieved once) and also about the state of the environment (e.g., whether a task can be allocated at a current state q_1 (today) or at potential states that follow q_1 (tomorrow)). Then it is reasonable to allocate a task to who—either a single agent or a group of agents—is *capable* of fulfilling it. While some consider this as strategic ability, we emphasize the importance of knowledge in “*being capable of doing something*”. Basically, as highlighted by [1], the strategic ability of an agent group is limited to their knowledge of the system and its dynamics (e.g., in a physical confrontation, even if agent a_i is strong enough to capture an adversary agent a_j , not knowing that a_j is located in front of a_i , avoids it exercising its potential). This is why we focus on *uniform* strategies (see Section 2.2).

Prior to allocating tasks and building on the notion of a uniform strategy, we say a local state of affairs G_q is *suitable* for a multiagent system (the *suitability* principle) only if the grand coalition Σ has a uniform strategy in state q to ensure it (note that we are referring to components of a particular CEGS \mathcal{M} that models a given MAS). Then given a suitable G_q , a task allocation is *valid* (the *validity* principle) only if we assume the compliance of agents with allocated tasks entail that G_q . We later discuss the third principle on *justifiability* of task responsibility. In the following, we specify the task allocation component of TasCore and show its desirable properties.

Definition 1. *Given a multiagent system \mathcal{M} , a local state of affairs $G_q = \varphi_1 \wedge \dots \wedge \varphi_m$, and the assignment of φ_i ($1 \leq i \leq m$) to agent group $\Gamma_i \subseteq \Sigma$, the assignment set $\{\langle \varphi_1, \Gamma_1 \rangle, \dots, \langle \varphi_m, \Gamma_m \rangle\}$ is a TasCore task allocation iff (1) Γ_i is a minimal group with a uniform strategy in q to ensure φ_i and (2) for any two intersecting groups Γ_i and Γ_j ($1 \leq j \leq m$), $\Gamma_i \cup \Gamma_j$ is a minimal group with a uniform strategy in q to ensure $\varphi_i \wedge \varphi_j$.*

⁵ “*Who gets what, and why?*” is the focus of the next section (we acknowledge the title of [33]).

As discussed earlier, this approach for allocating tasks captures the strategic abilities of agents and their epistemic limitations for allocating tasks. (We highlight that the allocation process is based on perfect information about agents, their abilities, and knowledge. Our reference to imperfect information is to the information that is available to involved agents in the MAS.) The following theorem shows that for a suitable state of affairs, there exists a task allocation that satisfies the two conditions in Definition 1. In q , we modelcheck to find minimal agent groups capable of ensuring propositional components of G_q and generate a task allocation that gives the task of ensuring each component φ_i (the i th components of G_q) to an agent group Γ_i which has a uniform strategy to ensure φ_i . This procedure generates a valid allocation if G_q is a suitable state of affairs.

Theorem 1. *Given a suitable G_q , there exists a valid task allocation in q .*

Proof. We provide a constructive proof by presenting a task allocation procedure based on ATL_{ir} model checking. First, for all $\varphi_i \in G_q = \{\varphi_1 \wedge \dots \wedge \varphi_m\}$, we use standard ATL_{ir} model checking [23] and apply a minimality-checking loop to generate the set, denoted by Φ_i , of minimal agent groups $\Gamma \subseteq \Sigma$ with the ability to ensure φ_i from q . Given all Φ_i , the set of allocation tuples $\{\langle \varphi_1, \Gamma_1 \rangle, \dots, \langle \varphi_m, \Gamma_m \rangle\}$ in which the two conditions of Definition 1 are satisfied is non-empty thanks to the suitability of G_q . ■

Note that we take the collective state of affairs and allocate each component to capable groups—with a uniform strategy to fulfill it. Recalling the notion of uniform strategy (in Section 2.1), given that the group has a collective strategy, each agent has an individual strategy that contributes to the fulfillment of the task. This way of allocating tasks to group, leaves it to the group to decide what ultimate individual action (in each state) individuals should take to see to it that the task is fulfilled. This is to see agents in an organizational setting as autonomous entities to whom we do not tell what exact action to take. We see each agent as a group member able to collaborate to execute a strategy, based on a repository of actions, such that the task that is allocated to the group is fulfilled. We later show how we can (retrospectively) ascribe a degree of responsibility to each individual based on what they did (outcome of collective actions) and what they had to do (allocated tasks).⁶

⁶ One may argue that task allocation is unnecessary as we can simply allocate G_q to Σ . While such a suggestion works under the perfect information assumption, the epistemic limitations of agents means such a simplistic approach will not work under

As discussed, the notion of the *collective state of affairs* G_q is a local notion to specify what properties are valuable in an organizational setting. Hence they ought to be satisfied by agents in Σ collectively. G_q says what is expected to be satisfied being in state q assuming that it contains expectations that were given previously and are not yet satisfied but valuable. In other words, agents are not required to keep a repository of tasks. This enables the expression of real-life situations where dynamic task allocation is desirable (i.e., it allows changing the state of affairs as the system evolves, hence gives the ability to give a different task to a group). Our approach to consider a local state of affairs gives a form of deterministic Markov property [27] to task allocation in TASCORE. In other words, local suitability can be extended and if satisfied globally (in all states) guarantees the existence of a valid task allocation in all states of MAS.

Proposition 1. *If G_q is suitable for all states $q \in Q$, the procedure presented in (the proof of) Theorem 1 generates a valid task allocation in all states regardless of the evolution of the system, represented by a materialized q -history $h \in \mathcal{H}$.*

Proof. Note that G_q is a local notion (on state q) and that its suitability is independent of any q -history h . Then to prove, we need to show that having a G_q suitable in all $q \in Q$, there exists a valid task allocation in each q . This is given by Theorem 1. ■

Under TASCORE, agent groups to which we allocate a task may intersect. E.g., if we allocate φ_1 to Γ_1 and φ_2 to Γ_2 , agents in $\Gamma_1 \cap \Gamma_2$ ought to take a part in ensuring both φ_1 and φ_2 . Then one may ask whether in such cases, agents in the intersection are supposed to choose between two alternatives and only satisfy one task while dismissing the other one. The following proposition shows that using the proposed procedure in Theorem 1, allocated tasks are not mutually exclusive.

Proposition 2. *Given a suitable G_q , there are no two exclusively satisfiable tasks allocated to a group.*

Proof. Follows directly from the second condition (in Definition 1) that any generated allocation—by the procedure presented in Theorem 1—satisfies. Note that the concern is not only about mutual exclusivity in a *logical* sense, but in a

imperfect information. This is because a uniform strategy for Γ to ensure a property is not necessarily uniform for all super-groups $\Gamma' \supset \Gamma$ as knowledge does not necessarily grow monotonically in groups.

strategic sense (i.e., that the available strategies for a group to fulfill two properties are coherent). ■

In organizational settings, agents are assumed to be a part of a collaborative practice, hence ought to fulfill their set of allocated tasks. But a valid question is whether they have any rational reason to avoid fulfilling and deviate from their tasks. In other words, whether an allocation is stable (in game-theoretic terminology). The following theorem shows that a TasCore task allocation is stable as no group has a rational incentive for deviation.

Theorem 2. *Given a suitable collective state of affairs G_q , a TasCore task allocation under the procedure presented in Theorem 1 is stable in the sense that no agent group has a rational incentive to deviate from its allocated task(s).*

Proof. Building on Proposition 2, for any allocated task to a group Γ —under the proposed procedure—any group Γ has a uniform strategy to fulfill the task. Hence, no group in Σ has a rational incentive to deviate from the allocation (i.e., to not fulfill the tasks). Note the assumption that agents are a member of the organization, hence prefer to deliver their tasks if they are able to do so. ■

Note that we are focused on the availability of a “task-fulfilling” strategy to agent groups—assuming that allocated tasks ought to be fulfilled if such a strategy exists—and abstract from agents’ preferences.

Example 1. In our example (see Figure 1), there are various ways (with different levels of control) which the parents can use to express their expectation. One way would be to have $G_{q_0} = \{\circ f, \circ\circ s, \circ\circ\circ t\}$ which is an explicit form with temporal limitations. Another more-relaxed form is $G_q = \{\circ f, \diamond s, \diamond t\}$. For both forms the first element can be allocated as a task to $\{A, B\}$, the second one to $\{C\}$, and the last one to $\{A, C\}$. Note that for a strategy to be uniform for a group, it should be accessible to them from all the indistinguishable states. In this case: $\{A, B\}$ has a strategy to ensure $\circ f$ from both q_0 and q_1 ; $\{C\}$ can ensure s either eventually ($\diamond s$) or in the state after the next ($\circ\circ s$); and $\{A, B\}$ can ensure tidiness.

5 Ascribing Responsibility in TasCore

Assuming that the allocation of tasks to agents definitely results in the fulfillment of tasks—and in turn brings about the collective state of affairs—is unreasonable in real-life environments. This is because autonomous agents are

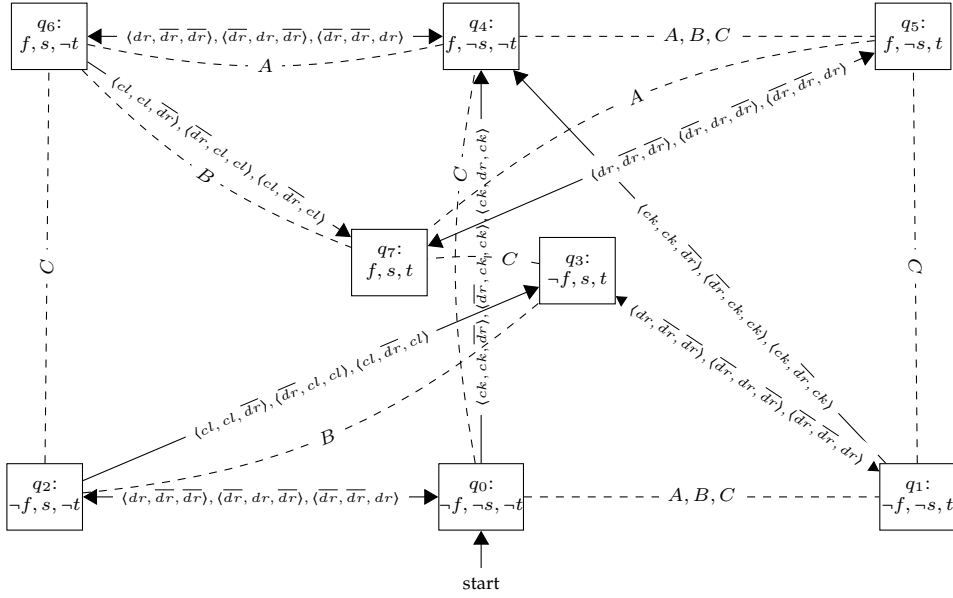


Fig. 1. Set of agents $\Sigma = \{A, B, C\}$ (Alex, Bob, Cathy), set of actions $Act = \{ck, dr, cl, id\}$ (cooking, driving, cleaning, idle, and for any $act \in Act$, \overline{act} denotes any action in $Act \setminus \{act\}$), and set of propositions $\Pi = \{f, s, t\}$ (ready food, being at spot, tidiness). For any unmentioned action profile α and arbitrary $q \in Q$, we have that $o(q, \alpha) = q$ (i.e., we avoided an α self-loop on every node). To model epistemic limitations, we assume that cooking (act ck) is only possible at home while cleaning (act cl) is only possible at the picnic spot. To model epistemic limitations, some states are indistinguishable for some agents represented by dashed lines labeled with the agent(s) who can not distinguish the states that the dashed line connects.

not artifacts but entities that may opt to exercise their autonomy and do otherwise. In organizational settings, such undesirable behavior is possible. But in addition—and based on the assumption that agents are a member of the organization, hence expected to deliver their allocated tasks—we can ascribe a degree of responsibility to agents. In particular, to those who contributed to a collective state of affairs being unfulfilled. This form of reasoning is known in the literature on responsibility as *backward-looking responsibility* [31, 36]. Basically, the reasoner observes the materialized history of events/actions (that lead to a given outcome situation S) and with the knowledge about agents' available actions, ascribes responsibility to agents who contributed to the occurrence of S or to those who could avoid $\neg S$ but apparently did not exercise their avoidance potential. The former causality-oriented approach is known as *causal responsibility* [12] based on the potential to bring about, while the latter is known as *strategic responsibility* [36] based on the potential to avoid. One may refer to this backward-looking form of responsibility reasoning as evaluating the blameworthiness if S is known to be a normatively undesirable state of affairs.

In our picnic scenario, imagine that the whole group arrives at the picnic spot with no food (states q_2 or q_3 in Figure 1). Who is responsible for such an undesirable outcome? Which agent(s) or agent group(s) can be blamed? And to what extent? One can look at the history of events (i.e., who did what at any state in the chain of states that ends in the current state), check the list of allocated tasks in those states, and justifiably blame a group of agents that ought to deliver the task of preparing food. This procedure is in particular a justifiable one in imperfect information settings if the agents' strategic ability and their epistemic limitations were taken into account in the preceding task allocation procedure. To ascribe responsibility to agents, we adopt the approach described in [36]. The responsibility reasoning notion of [36] is in-line with our approach in TasCore as we both focus on imperfect information settings. Following this, we see a group Γ as being responsible for an outcome φ in state q given a history h if φ holds in q while ensuring $\neg\varphi$ was among the allocated tasks to Γ in a state (other than q) in h . Formally:

Definition 2. *Given a multiagent system \mathcal{M} , TasCore task allocation, and the materialized q -history h , an agent group Γ is q -responsible for φ iff (1) φ does not hold in state q and (2) φ is among the allocated tasks to Γ in a state $q' \in h$ for $q \neq q'$.*

In this view, a group that is tasked to ensure φ is responsible for it until φ becomes true. Note that being responsible for φ does not imply any negative connotation here. (That is why we avoid using the negatively-loaded term “blameworthy” here.) This is crucial to note because a group might have the plan to bring about φ in a next state or maybe they are assigned a new task in the dynamic setting of TasCore. Moreover, note that a φ being unfulfilled in q directly implies that any state of affairs that contains φ is not satisfied. For instance, in our picnic scenario, having no food in q_2 implies that the collective state of affairs “well-organized picnic” $G_{q_0} = \{f, s, t\}$ that was given to the multiagent system in q_0 is unsatisfied in q_2 . This notion of responsibility can be seen as a measure of the collective state of affairs satisfaction in organizational settings. If ensuring G_q is given to a MAS, a state in which no group is responsible for any $\varphi \in G_q$ is a state where G_q is fulfilled completely.

Proposition 3. *Given a suitable state of affairs G_q , the TasCore task allocation, and the materialized history $h = q, \dots, q'$, we have that G_q is satisfied in q' if no agent group Γ is q' -responsible for any $\varphi \in G_q$.*

Proof. Having no group Γ being q' -responsible for any $\varphi \in G_q$ implies that for all φ , either of the two conditions in Definition 2 does not hold in q' . As TasCore task allocation is applied, the second condition holds for all φ . Thus the first condition does not hold for all φ , i.e., that all the elements of G_q are satisfied in q' . ■

Theorem 3. *Given a suitable G_q , there exists a nonempty set of states $S \subseteq Q$ such that for all $q' \in S$, no group Γ is responsible for any arbitrary $\varphi \in G_q$ based on history q, \dots, q' .*

Proof. In q , such a set in which $\bigwedge_{\varphi \in G_q} \varphi$ holds is foreseeable as, for ensuring a suitable G_q , a uniform strategy is available to Σ . ■

As we discussed, responsibility ascription is to agent *groups*. Then a problem that we face in multiagent systems—known as the *responsibility gap* problem—concerns situations where a group of agents is responsible for an outcome but the share of each individual agent is not clear [28]. In principle, the question is about the extent of responsibility of each agent for an unfulfilled task. This is a translation from group responsibility to individual responsibility. We highlighted that as we are dealing with tasks in TasCore, responsibility sharing techniques for MAS (e.g., [36, 12]) are not directly applicable. In TasCore, an

agent might be in two different groups with different tasks (coalitional dynamics) and moreover, may get involved in different tasks as the system evolves (temporal dynamics). Note that if agent group Γ with n members is found to be responsible for a given φ , sharing the responsibility equally is not a reasonable approach as each individual possesses different levels of knowledge and ability, and hence had different levels of contribution to φ . A standard approach is to consider *fairness* properties⁷:

Theorem 4. *Given a suitable collective state of affairs G_q , TAsCore task allocation, and q' -history $h = q, \dots, q'$, the degree of q' -responsibility of agent $a \in \Sigma$ for $\varphi \in G_q$, denoted by $\mathfrak{R}(a, \varphi, h)$, satisfies the fairness properties if $\mathfrak{R}(a, \varphi, h) = \Phi(a, \rho)$ where Φ returns the Shapley value of agents in the cooperative game $\langle \Sigma, \rho \rangle$ in which $\rho(\Gamma \subseteq \Sigma)$ is equal to 1 if $\Gamma' \subseteq \Gamma$ is q' -responsible for φ based on h and 0 otherwise.*

Proof. The game is such that—using the Shapley value—each agent receives a degree of responsibility with respect to its contribution to responsible groups. ■

Thanks to the properties of the Shapley value, this degree of responsibility is a fair way for bridging the responsibility gap and distributing the responsibility of each individual in TAsCore. As a direct result we have:

Proposition 4. *For a given $\varphi \in G_q$ and history h , we have that $\sum_{a \in \Sigma} \mathfrak{R}(a, \varphi, h) \in \{0, 1\}$.*

Proof. Having h , either q is a state in h or not. If not, then φ is not among the allocated tasks to any agent group in Σ hence the degree of responsibility for all agents $a \in \Sigma$ is zero. The same holds if q is in h and φ holds in the last state of h . The only case in which some agent groups are responsible for φ based on h is when φ does not hold in the last state of h ; and in this case, the summation is equal to one thanks to the efficiency property of the degree. ■

Based on this result we have that in organizational settings, for any allocated but unfulfilled task, there exists a responsible group if the task manager follows TAsCore. Note that as we are assuming suitability (for the collective state of affairs) and discussing the ascription of responsibility degrees in the context of task coordination, impossibilities for which no group is responsible are avoided automatically. The next property is about the evolution of the notion of task responsibility and the degree of responsibility through a given history.

⁷ By fairness, we are referring to *Shapley*-based notion of fairness that possesses the properties of symmetry, additivity, efficiency, and dummy player [34].

Proposition 5. For a given $\varphi \in G_q$ and history h , the notion of responsibility and its degree are non-monotonic through the temporal evolution of h : formally, (1) being q_i -responsible for φ (for a $q_i \in h$) does not imply being q_{i+1} -responsible for φ . Moreover, (2) $\mathfrak{R}(a, \varphi, h'')$ (for $h'' = q, \dots, q', \dots, q''$) is not necessarily higher or lower than $\mathfrak{R}(a, \varphi, h')$ (for $h' = q, \dots, q'$ as a part the materialized history h'').

Proof. For part (1), we have that φ may hold in one state of the history but not in the next one. In the states where it holds, no group is responsible for it. For (2), we can rely on part (1) but also have that due to dynamism of TasCore , the task to satisfy φ can be given to new groups through h (e.g., to ensure a level of fault-tolerance). Therefore the number of responsible groups is not fixed temporally (through the history h). ■

Note that TasCore is neither aware of nor requires agents' preferences as it respects a *separation of concerns* in the process of responsibility ascription (i.e., it is not designed based on the knowledge about agents' internal settings, hence is operational in multiagent organizational settings under imperfect information). Moreover, we assume that being involved in TasCore implies that the agent is expected to dedicate its resources (represented in terms of its available actions in each state) to the organization and is expected to fulfill its allocated tasks.

Example 2. Having the allocation of tasks according to Example 1 and the history q_0, q_2, q_3 , the collective state of affairs G_{q_0} is not yet fulfilled. Going back through the steps of history, the task of preparing food was allocated to $\{A, B\}$ and they had a uniform strategy to avoid it remaining unfulfilled. In this case, they each have the degree of responsibility of $1/2$ due to their symmetric contribution.⁸

6 Discussion and Concluding Remarks

We discuss the relevance of task coordination under imperfect information, argue the implementability of TasCore by showing steps toward operationalization, and relate our contribution to past work.

Perfect vs. Imperfect Information Settings: Although assuming perfect information is realistic for closed environments (e.g., in production processes or

⁸ Note that in general, agents may have various forms of (asymmetric) contribution. Thus, it is not the case that the degree of responsibility of group members is equal in general (see [36]) for a discussion on various responsibility reasoning cases.

data-base systems), in most real-life applications, an agent’s knowledge of the environment, hence of the consequences of its acts, is limited. In our model, we allow the representation of agents with imperfect information and consider uniform collective strategies to capture the epistemic aspect of the notion of strategic ability. This means we can both allocate tasks in imperfect information settings and ascribe responsibility in a justifiable manner.

Note that TasCORE is also operational in perfect information settings—simply by assuming an empty indistinguishability relation (in $\text{CEGS } \mathcal{M}$) for all the agents in Σ . In this way, the presented complexity and experimental results in [16] for the perfect information scenarios can be applied to TasCORE if one intends to deploy it in a perfect information task coordination case. In our running example, if agents had perfect information, we could dismiss indistinguishability relations (i.e., by deleting the dashed lines in Figure 1). Then in the task allocation part of TasCORE , agents may have more strategies available to ensure the components of a given state of affairs—as they are not epistemically limited. This in turn affects the responsibility ascription process.

Implementability: As we presented our task coordination concepts and tools using concurrent game structures, then the logical characterization of TasCORE is standard. To that end, one can use the epistemic variant of ATL proposed in [22] that adds indistinguishably relations to explicitly specify the epistemic uncertainty of agents. This allows reasoning about the abilities and responsibilities of agent groups under imperfect information.

Providing such a logical characterization of our notions also enables the systematic verification of the two parts of TasCORE (for task and responsibility allocation). Building on this, one can use standard model-checking tools [25] to implement TasCORE as an operational task coordination tool and in turn enable its application in real-life problems (e.g., in the context of business management or collaborative industrial systems).

Given a state of affairs to be ensured by the MAS, TasCORE enables tasks to be allocated to agents or agent groups such that their fulfillment leads to the overarching state of affairs being ensured. For task allocation, TasCORE is the first model that considers both the epistemic limitations and strategic abilities of agents. Moreover, we allow the allocation of tasks not only to individuals but also to agent groups. We argue that allocation is not enough if we deal with autonomous agents. Thus, following the task allocation, one can use TasCORE to ascribe a degree of responsibility to agents with respect to a state of affairs.

This work is the first to employ the notion of strategic responsibility for task coordination in MAS.

Our approach to allocate tasks to agents is complementary to (*group*) *role assignment* in robot systems and open societies [37, 15, 17, 19]. While they generally assume that roles are to be taken by agents, in TAsCORE we allocate tasks to agents based on their ability to ensure a task and accordingly see them as being responsible for the outcome. In this way TAsCORE can be used in combination with multiagent organizational frameworks such as *Moise* [21] (more precisely, within each organizational unit and assuming the availability of knowledge about the unit to TAsCORE).

In relation to the multiagent responsibility and accountability literature [36, 35, 5, 12], our work applies responsibility reasoning (in specific, the ATL-based notion of strategic responsibility) for task coordination. Based on TAsCORE's degree of responsibility, one can ascribe blameworthiness and sanctioning penalties to agent groups (e.g., to enforce a given norm in MAS). We also highlight a related but distinguishable approach to this problem that is based on the notion of *causal responsibility* in [12, 18, 2]. Basically, causal responsibility (as presented in [12]) ascribes a degree of responsibility to agents based on their potential to *provide* a situation while strategic responsibility (as presented in [36]) ascribes responsibility based on their ability to *preclude*. We see these two perspectives as complementarily applicable in different domains.

In TAsCORE, we dismissed incentivization. We see that an interesting extension is to consider rewarding agents per task fulfillment to provide strategy-proofness. Otherwise, an agent may strategically block its own strategies to avoid the allocation of a task. Rewarding can encourage agents to employ their most "*effective*" strategy, and by effectiveness we are referring to a strategy that enables them to fulfill as many tasks as possible. In principle, rewarding agents following the fulfillment of a task (and sanctioning otherwise) nudges the behavior of economically-motivated and rational agents towards the fulfillment of collective-level organizational goals. To address this, we aim to integrate norm-aware *incentive engineering* techniques [6, 7, 14] into TAsCORE and consider *coalition forming* aspects [10, 20]. In such a line, the degree of responsibility can be used as a basis to formulate normative notions of blame/praiseworthiness which in turn enables developing sanctioning/rewarding mechanisms. Moreover, in TAsCORE, we merely focused on physical actions of agents. Thus, another extension is to enrich the repository of actions by adding an explicit repre-

sensation of communicative actions. Basically, through communicative actions, the agents' epistemic level may change. This extends `TasCore` and enables us to reason about subcontracting, delegation, commitment-based agreements, and in general scenarios in which agents have *mixed* strategies, consisting of physical and communicative actions.

References

1. Ågotnes, T., Goranko, V., Jamroga, W., Wooldridge, M.: Knowledge and ability. In: Handbook of Epistemic Logic. pp. 543–589 (2015)
2. Alechina, N., Halpern, J.Y., Logan, B.: Causality, responsibility and blame in team plans. In: AAMAS-2017. pp. 1091–1099 (2017)
3. Alechina, N., van der Hoek, W., Logan, B.: Fair allocation of group tasks according to social norms. In: CLIMA-2014. pp. 19–34 (2014)
4. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. *J. ACM* **49**(5), 672–713 (2002)
5. Baldoni, M., Baroglio, C., Micalizio, R.: Accountability, responsibility and robustness in agent organizations. In: Workshop on Responsible AI Agents. pp. 1–8 (2019)
6. Boella, G., Hulstijn, J., Van Der Torre, L.: Virtual organizations as normative multi-agent systems. In: HICSS-2005. pp. 192c–192c (2005)
7. Boissier, O., Gâteau, B.: Normative multi-agent organizations: Modeling, support and control. In: Dagstuhl Seminar Proceedings (2007)
8. Bulling, N., Jamroga, W.: Comparing variants of strategic ability: How uncertainty and memory influence general properties of games. *JAAMAS* **28**(3), 474–518 (2014)
9. Campbell, A., Wu, A.S.: Multi-agent role allocation: issues, approaches, and multiple perspectives. *JAAMAS* **22**(2), 317–355 (2011)
10. Chalkiadakis, G., Elkind, E., Markakis, E., Polukarov, M., Jennings, N.R.: Cooperative games with overlapping coalitions. *Journal of Artificial Intelligence Research* **39**, 179–216 (2010)
11. Chapman, A.C., Micillo, R.A., Kota, R., Jennings, N.R.: Decentralised dynamic task allocation: a practical game: theoretic approach. In: AAMAS-2009. pp. 915–922 (2009)
12. Chockler, H., Halpern, J.Y.: Responsibility and blame: A structural-model approach. *JAIR* **22**, 93–115 (2004)
13. Coutinho, L.R., Sichman, J.S., Boissier, O., et al.: Modeling organization in mas: a comparison of models. In: SEAS-2005. pp. 1–10 (2005)
14. Criado, N., Julián, V., Botti, V., Argente, E.: A norm-based organization management system. In: COIN-2009. pp. 19–35 (2009)
15. Dastani, M., Dignum, V., Dignum, F.: Role-assignment in open agent societies. In: AAMAS-2003. pp. 489–496 (2003)

16. Fatima, S.S., Wooldridge, M.: Adaptive task resources allocation in multi-agent systems. In: AGENTS-2001. pp. 537–544 (2001)
17. Frias-Martinez, V., Sklar, E., Parsons, S.: Exploring auction mechanisms for role assignment in teams of autonomous robots. In: Robot Soccer World Cup. pp. 532–539 (2004)
18. Friedenberg, M., Halpern, J.Y.: Blameworthiness in multi-agent settings. In: AAAI-2019 (2019)
19. Giordani, S., Lujak, M., Martinelli, F.: A distributed algorithm for the multi-robot task allocation problem. In: IEA/AIE-2010. pp. 721–730 (2010)
20. Habib, F.R., Polukarov, M., Gerding, E.H.: Optimising social welfare in multi-resource threshold task games. In: International Conference on Principles and Practice of Multi-Agent Systems. pp. 110–126. Springer (2017)
21. Hannoun, M., Boissier, O., Sichman, J.S., Sayettat, C.: Moise: An organizational model for multi-agent systems. In: Advances in Artificial Intelligence, pp. 156–165. Springer (2000)
22. Jamroga, W.: Some remarks on alternating temporal epistemic logic. In: FAMAS-2003). pp. 133–140 (2003)
23. Jamroga, W., Dix, J.: Model checking strategic abilities of agents under incomplete information. In: ICTCS-2005. pp. 295–308 (2005)
24. Kota, R., Gibbins, N., Jennings, N.R.: Decentralized approaches for self-adaptation in agent organizations. *TAAS* 7(1), 1:1–1:28 (2012)
25. Lomuscio, A., Qu, H., Raimondi, F.: Mcomas: A model checker for the verification of multi-agent systems. In: CAV-2009. pp. 682–688 (2009)
26. Macarthur, K.S., Stranders, R., Ramchurn, S., Jennings, N.: A distributed anytime algorithm for dynamic task allocation in multi-agent systems. In: AAAI-2011. pp. 701–706 (2011)
27. Markov, A.A.: The theory of algorithms. *Trudy Matematicheskogo Instituta Imeni VA Steklova* 42, 3–375 (1954)
28. Matthias, A.: The responsibility gap: Ascribing responsibility for the actions of learning automata. *Ethics and information technology* 6(3), 175–183 (2004)
29. Nair, R., Tambe, M., Marsella, S.: Team formation for reformation in multiagent domains like robocuprescue. In: Robot Soccer World Cup. pp. 150–161 (2002)
30. Pnueli, A.: The temporal logic of programs. In: SFCS-1977. pp. 46–57 (1977)
31. Van de Poel, I.: The relation between forward-looking and backward-looking responsibility. In: Moral responsibility, pp. 37–52. Springer (2011)
32. da Rocha Costa, A.C., Dimuro, G.P.: A minimal dynamical mas organization model. In: Handbook of Research on Multi-Agent Systems: Semantics and Dynamics of Organizational Models, pp. 419–445. IGI Global (2009)
33. Roth, A.E.: *Who Gets What, and Why: The New Economics of Matchmaking and Market Design*. Houghton Mifflin Harcourt (2015)

34. Shapley, L.S.: A value for n-person games. *Contributions to the Theory of Games* 2(28), 307–317 (1953)
35. Yazdanpanah, V., Dastani, M.: Distant group responsibility in multi-agent systems. In: *PRIMA-2016*. pp. 261–278 (2016)
36. Yazdanpanah, V., Dastani, M., Jamroga, W., Alechina, N., Logan, B.: Strategic responsibility under imperfect information. In: *AAMAS-2019*. pp. 592–600 (2019)
37. Zhu, H., Zhou, M., Alkins, R.: Group role assignment via a kuhn–munkres algorithm-based solution. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 42(3), 739–750 (2011)