Towards Team Formation via Automated Planning

Christian Muise*, Frank Dignum†, Paolo Felli*,
Tim Miller*, Adrian R. Pearce*, Liz Sonenberg*

*Department of Computing and Information Systems, University of Melbourne
{christian.muise, paolo.felli,tmiller, adrianrp, l.sonenberg} @ unimelb.edu.au

†Department of Information and Computing Sciences, Universiteit Utrecht
f.p.m.dignum@uu.nl

Abstract. Cooperative problem solving involves four key phases: (1) finding potential members to form a team, (2) forming the team, (3) formulating a plan for the team, and (4) executing the plan. We extend recent work on multi-agent epistemic planning and apply it to the problem of team formation in a blocksworld scenario. We provide an encoding of the first three phases of team formation from the perspective of an initiator, and show how automated planning efficiently yields conditional plans that guarantee certain collective intentions will be achieved. The expressiveness of the epistemic planning formalism, which supports modelling with the nested beliefs of agents, opens the prospect of broad applicability to the operationalisation of collective intention.

1 Introduction

It is both a challenging and important problem to form a cohesive team that can achieve a task. Wooldridge et al. [24] propose four key phases to cooperative problem solving: (1) *potential recognition* where the team "initiator" must identify the capabilities of the agents; (2) *team formation* where the potential team members are persuaded to join for the collective intention; (3) *plan formation* where a plan is constructed; and (4) *plan action* where the joint plan is executed. Dignum et al. [5] propose a framework for these four stages that relies on structured dialogue between the initiator and the agents in the domain. In this work we introduce a novel approach for building principled and scalable mechanisms for team formation that exploits recent advances in multi-agent epistemic planning, and we illustrate the approach by working with a model of team formation inspired by the Dignum et al. framework (referred to as DDV in the remainder of the paper).

We focus on the *initiator* role in team formation, which involves assessing the potential for team formation and persuading possible members to join. While all four phases of team formation are important, currently we address the first three phases only: i.e., the phases that involve the initiator's deliberation prior to the execution of a plan. We take the view that the initiator will deliberate about what questions to ask, what promises to make, the composition of the team, and the potential of achieving the overall objective as a team, *all prior to the execution of any action or dialogue*. While the initiator is not strictly required to form a plan before the actual dialogue occurs, doing so can save the initiator from asking irrelevant questions and performing actions that ultimately will never lead to a solution. This is vital particularly when part of the dialogue may involve making promises to the agents that make up a team as part of the persuasion. Thus, the initiator not only plans for the required dialogue, but also for the eventual plans contingent on the possible dialogue outcomes.

The four phases of teamwork formation are general concepts. Here, we consider a specific realization of the four phases where the initiator can: (1) ask agents about their "capabilities" (e.g., can agent 4 lift blue blocks?); (2) ask and convince agents about their "intentions" to assist in a task given a particular promise (e.g., will agent 3 lift red blocks if we promise to put block 4 in room 2?); and (3) orchestrate the actions of the agents that agree to assist. The reasoning task for the initiator is to come up with a conditional plan (conditioned on the responses of the agents) such that a cohesive team can be formed to achieve the overall objective. This cannot always be guaranteed (e.g., if every agent refuses to help), but the initiator's deliberation process should at least discover the ways in which a successful team can be formed.

We model the problem from the perspective of the team initiator using an extension of the recently introduced multi-agent epistemic planning (MEP) formalism [16], which uses syntactic belief bases restricted to non-disjunctive clauses to represent nested agent beliefs [14]. MEP extends classical planning by allowing the nested belief of agents in a multi-agent environment as action preconditions and effects, and nested beliefs as goals that can be posed. Using MEP allows us to model the critical notion of an agent's belief that their own objectives have been satisfied. We extend MEP by allowing for non-deterministic action outcomes – providing a natural way to express yes / no questions for dialogue – and by using a generic fragment for the action theory that allows for team formation to take place. All of the existing MEP domain descriptions, which describe the actions that agents can take in the domain and the effect that they have on the belief of agents, can easily be plugged into the augmented system.

The realization of our approach using automated planning is both powerful and flexible. Unlike other approaches, such as BDI [19] or hierarchical plan representations [9], the plans for the agents need not be specified in advance. Rather, we can use the powerful automated planning techniques that have been developed over the recent decades to synthesize the viable plans for us [10]. This approach shifts our focus from one of creating a new solving technique to one of creating a novel encoding for existing solvers.

For our running example, we adapt the Blocks World for Teams (BW4T) domain [11] to include agents with varied capabilities and tasks that can require multiple agents, a natural extension given our focus on modelling team formation. In BW4T, agents carry different-coloured blocks around various rooms. In our adaptation, agents have capabilities to carry only certain colours (e.g., blocks with the blue colour can only be lifted by "blue lifters"), blocks may take on multiple colours, and agents may have multiple colour capabilities.

In the next section we provide the necessary background and notation for our approach. Following this, we describe how we have modelled the problem of teamwork formation in Section 3 and encoded it for automated planning in Section 4. We then discuss a preliminary evaluation using the BW4T domain in Section 5 and conclude with a discussion of related and future work in Section 6.

2 Background

2.1 Team formation

Team formation based on various approaches to 'matching' potential participant skills with the requirements of a task have long been studied. Some approaches involve heuristics guided by logical analysis, e.g. [12, 23], others involve formal mechanisms based on multimodal logics, e.g. [6, 7], and others draw on game theoretic and optimisation techniques, e.g. [1, 2, 4, 13, 18, 20].

Although all this work is about team formation different approaches tend to focus on different parts of the issue of forming a team to accomplish a task. In the optimisation work the focus is generally on finding the best team given that it is clear which plan (or set of tasks) is to be executed [4, 18]. Thus the question is how to allocate tasks to agents in an optimal way. In work related to coalition formation, e.g. [21], the emphasis is typically on the negotiation process between the agents in order to join and stay within a team. This can be done using game theoretic notions, in which division of possible rewards over a group play an important role. It can also be done using argumentation in which the emphasis shifts to the reasons for joining a team and persuading potential team members about the justification or importance of the team goal and or a particular plan to reach the goal, e.g. [2]. In the logic based approaches the emphasis is on the exchange of information about goals, intentions and beliefs such that the logical pre-conditions for working as a team according to the SharedPlans framework [8] are fulfilled.

In DDV the emphasis is on what is needed for a set of agents to start working as a team to achieve a joint goal. It involves at least that all the agents agree upon their role in the plan to achieve that goal (or in other words the tasks that they are willing to perform within the plan) and that they have enough information to execute their task at the right moment in time. As mentioned, we consider team formation from the perspective of the *initiator*. The first task of the initiator is to form a partial (abstract) plan for the achievement of the (team) goal. On the basis of the (type of) subgoals that it recognizes, it will determine which agents might be most suited to form the team. In order to determine this match, the initiator seeks to find out the properties of the agents, with the DDV framework focusing specifically on three aspects: their *abilities*, *opportunities*, and *willingness* to participate in team formation. Ability does not depend on the situation, but is taken as an inherent property of the agent. The aspect of opportunity takes into account the possibilities of task performance in the particular situation, involving resources and possibly other properties. The aspect of willingness considers the agents' mental attitudes towards participating in the proposed team goal. The outcome of the potential recognition stage is that the initiator knows whether or not it is possible to form a team, but has yet to engage in *team persuasion*, i.e. persuading potential team members to take on the intention to achieve the overall goal.

As our focus is on the initiator's reasoning process, we adopt a slightly altered view of the notions "goal" and "collective intention". For our work, the initiator's *original goal* is the specification of what the initiator would like to achieve as a result of forming and directing a team. The *collective intention* of this team will include both the team initiator's original goal and any subgoal arising from the team formation process.

In the rest of this paper we will show how multi-agent epistemic planning can be used to operationalize this approach, and provide a practical way to generate possible plans for the team to achieve a goal.

2.2 Multi-agent Epistemic Planning

We adopt a formalism of planning where the planning agent can reason in a limited fashion about the nested belief of other agents in the domain [16]. The state of the world in this setting is a collection of *Restricted Modal Literals* (RMLs) which are taken from the set $\mathcal{L}_{\mathcal{F}}^{Ag,d}$ defined by the following grammar:

$$\phi ::= p \mid B_i\phi \mid \neg\phi$$

where p is from a set of primitive fluents $\mathcal{F}$ and i is from a set of agents Ag. The modal proposition $B_i\phi$ states that agent i *believes* proposition ϕ, in which ϕ can be other possibly-nested beliefs. The maximum depth of nesting is limited by d. If $\mathcal{F}$, Ag, and d are all finite, then so is the set of RMLs $\mathcal{L}_{\mathcal{F}}^{Ag,d}$. Following earlier work [16], we define a *Multi-agent Epistemic Planning* (MEP) problem as the tuple $\langle \mathcal{F}, \mathcal{A}, \mathcal{I}, \mathcal{G}, Ag, d \rangle$, where:

- $\mathcal{F}$ is a set of atomic fluents.
- $\mathcal{A}$ is a set of actions (described below).
- $\mathcal{I}$ is a subset of $\mathcal{L}_{\mathcal{F}}^{Ag,d}$ describing the initial state.
- $\mathcal{G}$ is a subset of $\mathcal{L}_{\mathcal{F}}^{Ag,d}$ describing the goal condition.
- Ag is the set of agents in the domain.
- d is the maximum depth of nesting allowed.

For every action a in $\mathcal{A}$, we will use $\textsc{Name}(a)$ to indicate the action's name. $\textsc{Precond}(a)$ is the subset of $\mathcal{L}_{\mathcal{F}}^{Ag,d}$ that must hold in order for a to be executable, and $\textsc{Effects}(a)$ is a *set of* outcomes, of which exactly one will occur after a is executed: i.e., the action outcomes may be *non-deterministic* [3]. The possible outcomes of an action are known in advance, but the precise outcome is known only after the action has been executed. Thus, we are assuming fully-observable, non-deterministic (FOND) planning, in which actions are non-deterministic, but their effects are fully observable after execution. This is in contrast with the original MEP formalism where every action was necessarily deterministic.

This generalization of deterministic actions is an appealing way to model dialogue. The modifications we made to accommodate for non-deterministic actions did not change the theoretical framework introduced by Muise et al. [16], as the non-determinism in the domain is fully observable (i.e., the agent will know which outcome occurs immediately after the action is executed). While this requires the planner to handle various contingencies depending on the action outcome, it does not alter the way beliefs are encoded using the standard MEP formalism. The only change made was to replace the classical sub-planner with a non-deterministic one. Using a non-deterministic planner allows us to plan for all contingencies offline, which can be extremely helpful in avoiding bad sequences of dialogue and bargaining actions.

Every outcome in $\textsc{Effects}(a)$ is a set of conditional effects that change the state of the "world", in which the "world" includes beliefs of agents. We use $cond \rightarrow f$ to signify the conditional effect that updates the state of the world for f to hold in the following state when *cond* holds in the current state. If the condition *cond* is empty, we will just omit the $\rightarrow$. If every action is deterministic, then a solution to a MEP problem is a sequence of actions that, when executed from the initial state, achieves the goal. As we allow for non-deterministic effects, a solution is generalized to be a policy mapping reachable states (including the initial state) to the action that the agent should execute next in that state.

Using a variation of MEP planning in lieu of classical planning provides valuable modelling properties in the context of multi-agent environments. We will point out some of these advantages throughout the paper in the context of our target domain, and refer the interested reader to [16] for a deeper discussion on how belief is maintained during the planning process.

2.3 Blocks World For Teams

As a testbed, we consider a modified version of the Blocks World For Teams (BW4T) domain [11]. In BW4T, agents must navigate a series of rooms to relocate blocks in a target goal configuration. We extend the general setting to include fluents indicating block types (each block can have one or more "colour" associated with it), as well as the capabilities for agents to lift blocks of a particular colour. The extension allows us to model the more complex setting of heterogeneous agents, as is typical with many team formation problems. The goal of the initiator will be to form a team that can collectively achieve the goal configuration of blocks. The task that we solve in this work is to synthesize a plan of dialogue steps that will yield such a team formation.

While not strictly required, we will associate an agent with every action that signifies the agent performing the action. For example, instead of the action *lift_blue_b1_room1* we will have actions *lift_i_blue_b1_room1* for every agent i in Ag (note that if block $b1$ was also of the *red* type, we would have a separate *lift_i_red_b1_room1* action for every agent). For simplicity we include the agent and objects in the action name, but in practice these are parameterized. To ensure only the appropriate agent lifts a block, preconditions will include the agent's capability: e.g., the *lift_ag1_blue_b1_room1* action will include *can_lift_ag1_blue* as a precondition (we discuss capabilities further in Section 3). Other preconditions and effects include the standard ones for the BW4T domain, as well as extra effects to update the belief of agents. The following is the full description for *lift_ag1_blue_b1_room1* (we have replaced repeated effects for each agent with a single effect for agent i):[1]

$$
\begin{aligned}
\textsc{Name}(a) &= \mathit{lift\_ag1\_blue\_b1\_room1}\\
\textsc{Precond}(a) &= \{\mathit{at\_ag}1\_\mathit{room}1, \mathit{block\_colour\_b}1\_\mathit{blue},\\
&\qquad \mathit{in\_b}1\_\mathit{room}1, \mathit{can\_lift\_ag1\_blue}\}\\
\textsc{Effects}(a) &= [\{\mathit{holding\_ag}1\_b1, \neg \mathit{in\_b}1\_\mathit{room}1,\\
&\qquad \mathit{at\_i\_room}1 \rightarrow B_i \mathit{holding\_ag}1\_b1,\\
&\qquad \mathit{at\_i\_room}1 \rightarrow B_i \neg \mathit{in\_b}1\_\mathit{room}1\}]
\end{aligned}
$$

[1] Note that, as with most blocksworld encodings, the fluent *in_b1_room* is false whenever an agent is holding block $b1$.

As a result of using the MEP framework, additional effects will be created to maintain certain properties. For example, the effect $at\_i\_room1 \rightarrow \neg B_i \neg holding\_ag1\_b1$ would be added to maintain consistency of belief; that is, if agent i believes that agent 1 is holding $b1$, then it cannot believe agent 1 is not holding $b1$. There are similar effects on the move actions; e.g. an agent will believe the contents of a room when they enter it. Using MEP gives us a much richer environment in which to pose our teamwork formation problem.

3 Model of Team Formation

The initiator must assess the capabilities of the agents, and bargain with them in order to convince them to join the team. As a side-effect of bargaining, the collective intention of the eventual team may change – every promise made during the bargaining phase will become a subgoal for the team to achieve in the final state of the plan in addition to the original goal.

As discussed earlier, we adopt a model of team formation inspired by Dignum et al. [5]. Because we wish to consider team formation from the perspective of the initiator, many of the details are abstracted away from the DDV model (e.g., the precise reasoning capabilities of the other agents). Further, from the perspective of an initiator that is considering the viability of forming a team, every agent is assumed to be "blindly committed": if they have agreed to join the team, they will perform the actions prescribed to them by the initiator as expected. Thus, where the DDV model concentrates on the formation of a joint intention for a team, this paper instead concentrates on the planning and willingness of other agents to participate in the plan. Note that we are not assuming that the agents are purely cooperative – the team initiator must *convince* them that it is worthwhile to join the team through a process of negotiation. In this section, we describe our model of team formation, and contrast it with that of DDV.

The objective of the initiator is to form a team that can achieve the initiator's original goal. As part of the team formation process, the initiator must ensure that the capabilities of the agents on the team will allow the goal to be achieved, and may also need to promise certain things for potential members to join the team. The problem that we address is how the initiator can reason about which questions to ask and bargains to make. Rather than isolating dialogue planning from reasoning about goal achievement, we model both concurrently. The advantage is that we can rule out certain bargaining options that provably will never result in a viable team.

The initiator agent takes the following steps to form a team that can execute a joint plan for the initiator's original goal:

(a) assess capabilities of the agents (Section 4.2);
(b) bargain with the agents about promises (Section 4.3);
(c) given the commitments made to the agents, plan for the collective intention of the team (Section 4.4).

Note again, that the initiator reasons about these steps before the communication actually takes place. Thus the plan is a conditional plan, based on the answers and commitments of the agents.

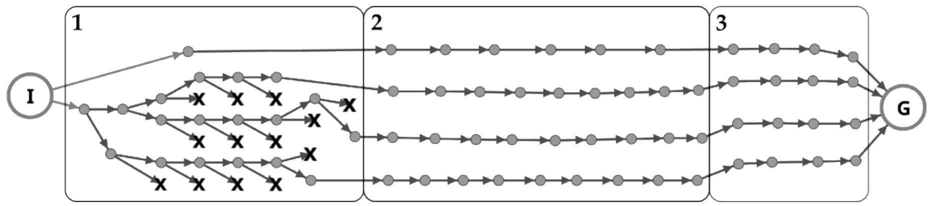

Fig. 1: Example plan for team formation. Stage 1: Capability assessment and bargaining. Stage 2: Planning steps to achieve the collective intention. Stage 3: Meta actions to ensure every promise was fulfilled.

Our teamwork formation model describes only the set of potential capabilities and bargains that the initiator should consider in the reasoning process – a process that takes place before the communication is initiated. Together with an action theory describing what the agents can do, an initial state of the world, and a goal configuration, the initiator synthesizes a conditional plan for forming a team to achieve the team's collective intention.

Thus the structure of a plan can be viewed as a tree of dialogue actions (branching occurring based on the agents' response to capability assessment or bargaining), where the leaves represent either a successful team's plan or a configuration of agent responses that lead to no viable team formation (e.g., if too many agents are unwilling). If there is a chance that making further bargains will allow a team to be formed, the planning phase will detect this. Naturally, the solutions produced will use only those agents necessary, and this is detected automatically from our encoding. Figure 1 shows a high-level structure of one such plan: the stage marked 1 is where the dialogue occurs, and the nodes marked 'x' are the situations when the dialogue fails (e.g., an essential agent refuses to help). We describe the other components in more detail below.

The key components of our team formation model are: (1) the potential capabilities of the agents; and (2) the range of bargains that can be offered to an agent. Formally, given the set of agents Ag and fluents describing the world $\mathcal{F}$, a teamwork formation model $\langle Cap, C, \mathcal{B}\rangle$ is a tuple where,

- $Cap \subseteq \mathcal{F}$ is a set of fluents representing the agents' capabilities;
- $C : Ag \to 2^{Cap}$ maps agents to the set of their potential capabilities; and
- $\mathcal{B} : Ag \to 2^{\mathcal{F}}$ maps agents to the set of the possible bargains they will consider.

Note that for agent $ag1$, $C(ag1)$ does *not* describe all of their capabilities. Rather, it describes a superset of the actual capabilities that agent $ag1$ might have. It is the responsibility of the initiator to surmise from agent $ag1$, which capabilities in the set $C(ag1)$ it actually has. Our notion of capability corresponds to the DDV notion of *ability*, and the initiator is able to ask the agents if they have particular capabilities as part of the information seeking dialogue.

After the initiator knows an agent is capable of what is required of them, the initiator must "persuade" the agent to join the team. The initiator achieves this through the use of *bargaining*. The initiator can offer any fluent in the set $\mathcal{B}(ag1)$ to agent $ag1$ as a promised subgoal that will hold at the *end* of the execution (i.e., in the final state of the system). $\mathcal{B}(ag1)$ may contain fluents involving other agents, so for example we may have a pair of agents that will only join the team if the other is promised a block: $\mathcal{B}(ag1) = \{holding\_ag2\_b1\}$ and $\mathcal{B}(ag2) = \{holding\_ag1\_b2\}$. This possibility of mutual support can lead to added restrictions on team formation, making the need even greater for the initiator to plan in advance.

After a bargain is proposed, the agent can agree or disagree, and the initiator plans for both eventualities. If the agent agrees, it becomes part of the formed team and achieving the promise must be satisfied by the plan. If the agent does not agree, then the initiator can try to persuade the agent in a different way, or try to find another agent to assist.

The plan must allow the promises to be satisfied in the following sense: if agent $ag1$ joined the team on the premise that $f \in \mathcal{F}$ is achieved, the goal of the initiator must now include $B_{ag1} f$. That is, a promise made to an agent must be believed by that agent when the plan's execution is complete. This allows for behaviour whereby the initiator can form a team where the members have inconsistent intentions, as long as the individual agents believe that their promises will be fulfilled in the end. We could keep the initiator "fully honest" by placing both f and $B_{ag1} f$ in the set of goals when a promise of f is made to agent $ag1$.

The set of promises that an initiator commits to, along with the original goal, then constitutes what we term the collective intention of the team. While the collective intention is not explicitly represented, as is the case with DDV, the plan produced by the initiator serves as an essential basis for the team to have collective intention. In a sense, the plan produced by our encoding is a certificate that the initiator can use in order to achieve DDV's form of collective intention during the actual dialogue phase.

The concepts of *willingness* and *team persuasion* from DDV are both covered by the initiator's ability to bargain. A key aspect of our approach is that the dialogue can occur in any order. The initiator can consider inquiring about capabilities, then bargaining with some agents, then inquiring about capabilities depending on the outcome of previous dialogue. This allows the initiator to condition their dialogue strategy based on the responses they have received so far.

The DDV notion of *opportunity* assumes that the initiator thinks they are both able and have the resources in the situation to achieve something. Instead of ascribing this notion to the agents, we task the initiator with assessing whether or not agents have the opportunity to achieve a subgoal. This is a natural consequence of our assumption that the initiator is reasoning about both the dialogue and planning phases simultaneously. This shifts the complexity of gauging an agent's opportunity to achieve a subgoal to the initiator's planning phase. The advantage of this is that the initiator does not need to reason apriori about which subgoals the agent has the opportunity to achieve (which can be a complex notion given the other agents that may be on the team). When the initiator considers a formed team, they can try to synthesize a plan with the team while *implicitly* computing the opportunities of every agent.

4 Encoding Team Formation

Our general approach to team formation is to model the cognitive process of a team initiator who must make various decisions about how best to form a quality team. The mental exercise of the team initiator involves not only the enumeration of team member configurations, but also the evaluation of a given team configuration's potential to achieve the goal. By considering everything from bargaining to physical actions in the world, the initiator can rule out bad team configurations and avoid unnecessary bargaining *before* launching the initial dialogue to form a team.

We begin by describing the general encoding we use, and then elaborate on the details particular to teamwork formation. Following the general encoding, we focus on our model of capability assessment and bargaining. These are the methods the initiator uses to make an informed judgement about who to include on the team. Next, we describe the three internal stages of the initiator's reasoning. These do not correspond directly to the phases of solving a joint task, but we do point out the relation between the two.

4.1 General FOND MEP Encoding

There are two sources of input for the generated FOND encoding: (1) the original MEP problem specification; and (2) the description of bargaining and capability properties for the agents that can form a team. Both will inform the fluents, initial state, goal configuration, and actions in the domain.

Fluents The MEP problem comes with a set of fluents $\mathcal{F}$ and agents Ag for the domain. These are combined, along with the maximum depth d, to generate fluents for the encoding that represent both what is true in the world and what the belief of each agent consists of. For example, the fluents in the BW4T domain FOND encoding will include $holding\_ag1\_b2$ and $Bag3\_in\_b1\_room4$.

Initial State The initial state will come directly from the MEP problem as well. By adopting the MEP framework, the initial state represents the belief of the initiator. It can either be fully specified (i.e., for every fluent f, either f or $\neg f$ holds in the initial state) or partially specified (i.e., the initiator is uncertain about certain facts). The option to use a partially specified initial state opens the door to a wider class of problems that include situations where the team initiator is *not* omniscient; a realistic assumption that often is overlooked.

Goal Configuration As with the initial state, we adopt the goal from the MEP problem for the FOND encoding. It will consist of a set of fluents that describe a partial state that must be achieved (e.g., having particular blocks placed in a specific location).

Actions The actions for the FOND encoding will correspond to those in the MEP action theory, with the precondition and effect RMLs replaced by their compiled fluent equivalent. We assume that every action from the MEP problem has an associated agent. For example, the action for picking up block $b1$ has a copy for every agent ($pickup\_ag1\_b1$, $pickup\_ag2\_b1$, etc). Similarly, the fluents that are required by an agent to conduct an action will have an agent associated with it (e.g., $hand\_free\_ag1 \in$ $\textsc{Precond}(pickup\_ag1\_b1)$).

4.2 Capability Assessment

To achieve the overall goal of the team, the initiator must assess the capabilities of potential members. Some capabilities may be known in advance, but in general we assume that the initiator must consider "asking" the agents if they are capable of certain tasks. For example, the initiator may ask agent 1 if he or she can lift blue blocks. Because the initiator deliberates offline, she must consider all of the possible outcomes of a question. For the time being, we limit the form of the question to simple yes / no inquiries such as the example above. After asking the question, the initiator will continue the deliberation process in two ways: once assuming a positive response, and again assuming a negative response. Thus, we can model the question using a non-deterministic action:

$$\begin{aligned}\textsc{Name}(a) &= \mathit{ask\_if\_ag1\_can\_lift\_blue\_blocks}\\ \textsc{Precond}(a) &= \{\neg \mathit{can\_lift\_ag1\_blue}, \neg \mathit{cannot\_lift\_ag1\_blue}\}\\ \textsc{Effects}(a) &= [\{\mathit{can\_lift\_ag1\_blue}\}, \{\mathit{cannot\_lift\_ag1\_blue}\}]\end{aligned}$$

There are two important aspects of this encoding. First, one outcome will allow the initiator to orchestrate standard actions for lifting blue blocks using agent 1 (recall that a precondition of $\mathit{lift\_ag1\_blue\_b2\_room1}$ is that $\mathit{can\_lift\_ag1\_blue}$). Second, both outcomes make it impossible to ask this question a second time. This second aspect is important because we do not want to assume that repeating a question will eventually lead to a different response.

4.3 Bargaining

Once the initiator is confident that an agent has the right capability for the task, she must ensure that the agent is willing to help. This notion corresponds directly to the idea of persuasion in DDV. Rather than convincing the agent that the collective intention is achievable through dialogue, the initiator will consider making "promises" about the collective intention of the team. For example, she might tell agent 2 that if they decide to join the team, they can have block 3 at the end of the sequence.

The combination of the original goal with the set of promises made to the agents constitutes the collective intention for the team. Every agent on the team is either willing to join for free, or is willing to join for a particular price. We model this aspect of the dialogue in a fashion similar to capability assessment using the following non-deterministic action:

$$\begin{aligned}\textsc{Name}(a) &= bargain\_with\_ag1\_for\_holding\_ag1\_b2\\ \textsc{Precond}(a) &= \{\neg ag1\_willing, \neg ag1\_unwilling\_holding\_ag1\_b2\}\\ \textsc{Effects}(a) &= [\{promised\_ag1\_holding\_ag1\_b2, ag1\_willing\},\\ &\qquad \{ag1\_unwilling\_holding\_ag1\_b2\}]\end{aligned}$$

Similar to the capability assessment, the initiator cannot try continually to bargain using the same offer. However, different actions may correspond to different promises that the initiator can propose in order to convince the agent to join the team. As mentioned previously, an agent can only perform an action if they are willing to be a part of the team (i.e., $ag_i\_willing$ holds). The additional fluent $promised\_ag1\_holding\_ag1\_b2$ maintains this aspect of the collective intention, and we will see next how the reasoning for the initiator ensures that the team achieves this subgoal.

4.4 Three Stages for Initiator Reasoning

The initiator's deliberation process is encoded as a FOND MEP problem, and will go through three distinct stages: (1) forming the team based on capability assessment and bargaining; (2) constructing the plan for the team; and (3) ensuring that the goal is satisfied and the promises fulfilled. In relation to the four phases of a team jointly achieving a task, stage (1) corresponds both to team assessment and team formation (phases 1 and 2 in cooperative problem solving [24]) – the initiator has the option to consider interleaving the assessment of capabilities and persuasion. Stages (2) and (3) correspond directly to plan formation (phase 3). Finally, as mentioned earlier, we do not consider the fourth phase from Wooldridge and Jennings that covers plan execution. We have marked the actions belonging to the three stages in an example plan shown in Figure 1.

To restrict the reasoning to each of the stages, we include the following components in the encoding:

(a) We introduce the auxillary fluents $stage\_formation$, $stage\_planning$, and $stage\_final$, with $stage\_formation$ set to true in the initial state.
(b) We introduce two actions, $start\_planning$ and $finish\_planning$, that simply change the value of the auxillary fluents appropriately (i.e., removing the current stage and adding the next).
(c) Every capability assessment and bargaining action has the extra precondition of $stage\_formation$.
(d) Every standard action has the extra precondition of $stage\_planning$.
(e) We introduce new "satisfy" actions for each agent, as described below.

The above changes force all of the capability assessment and bargaining to occur before the standard planning actions are considered. Doing so forces the initiator to consider forming the team prior to considering if they can achieve the goal. However, note that the initiator reasons about all three stages *before* they physically start any dialogue.

For the final phase (i.e., stage (3)), the initiator uses the plan from stage (2) to determine if the team has achieved the collective intention. The initiator must achieve both the original goal and the presumed satisfaction of every agent. The initiator will presume an agent satisfied if either the agent was not part of the team to begin with, or by acknowledging that the agent believes any promise made to them was kept. The first of the two actions, which covers the case of an agent not part of the team, is as follows:

$$\begin{aligned}
\textsc{Name}(a) &= \textit{satisfy_ag1_unwilling} \\
\textsc{Precond}(a) &= \{\neg ag1_willing, stage_final\} \\
\textsc{Effects}(a) &= [\{ag1_satisfied\}]
\end{aligned}$$

Note that we use $ag1\_willing$ here additionally to indicate that the agent was a part of the team, while $\neg ag1\_willing$ indicates that they were not. The other action capable of "satisfying" an agent is to reaffirm that they explicitly believe the promise that was presented to them during team formation:

$$\begin{aligned}
\textsc{Name}(a) &= \textit{satisfy_ag1_for_holding_ag1_b2} \\
\textsc{Precond}(a) &= \{stage_final, Bag2_holding_ag1_b2, \\
&\qquad promised_ag1_holding_ag1_b2\} \\
\textsc{Effects}(a) &= [\{ag1_satisfied\}]
\end{aligned}$$

The distinction of achieving an agent's bargained promise, as opposed to having that agent believe the promise is fulfilled, is an important one. On one hand, it is not enough for the team to achieve something that was promised to an agent while the agent remains unaware of this fact. On the other hand, this provides potential for deceitful behaviour – depending on how the agent updates their belief, they may *believe* that their objective holds when in fact it does not (e.g., they see a block placed in a room they desire, and then leave the room believing that it will remain there). This level of expressiveness is an intended consequence of using MEP planning as our underlying framework instead of generic automated planning techniques. Once all of the agents have been "satisfied", the reasoning process is complete. Note that because of the stage fluents, agents can be considered satisfied only in the final stage; during which the state of the world and the beliefs of the agents cannot be altered further.

It is worth emphasizing that the 3 stages are *all solved as a single planning problem*. The conceptual separation between the actions in each of the stages is implicit in any valid solution that the planner generates. Further, the planner is equipped with an efficient means of relevance analysis and policy reuse (see [17] for a discussion), which means that partial plans found for a subset of the agents can be reused in different configurations of the team. This reasoning essentially is "free" when we use a state-of-the-art FOND planner to solve the encoded problem. We would not get this benefit without significant overhead if we were to implement the three stages individually.

While perhaps counter-intuitive, the bundling of all stages into a single encoded problem allows us to fully leverage the planning technology at the core of our approach. Aside from the reuse of plan fragments mentioned earlier, the planner will also recognize when making a particularly bad decision early in stage (1) will prevent stages (2) or (3) to be successful. Further, we can use the produced conditional plan as a certificate for future negotiations by the team initiator (a phase out of scope for this work).

5 Preliminary Evaluation

We report on a preliminary evaluation to demonstrate the potential for solving teamwork formation problems with automated planning technology. We used the available implementation of the MEP framework [15], and wrote a compiler for team formation problems (cf. Section 3) that produces an encoded FOND MEP problem (cf. Section 4). To solve the encoded problems, we used an off-the-shelf FOND planner, PRP [17], which generates a policy for the initiator to follow. As we are using PRP in a black-box manner, we do not go into the details of how it computes a plan for the resulting encoding. The computed solutions may be suboptimal, but in general they do not contain superfluous dialogue or planning actions. By using modern planning technology, our approach is scalable to far larger problems; existing planners can solve problems with trillions of states in fractions of a second [10]. Where scalability suffers, as pointed out by Muise et al. [16], is when the depth of nesting and number of agents grows too large.

We modelled various settings for the BlocksWorld for Teams (BW4T) domain [11] with five agents, five rooms, four blocks, and three possible block colours, running four different scenarios. Each of the four scenarios tested a unique setting for the team formation (described further below), and details of the base encoding can be found in Section 2.3. Table 1 shows the time that it takes to synthesize a plan, including both the encoding and solving phases, as well as the final policy size measured as the number of possible states that it takes to encode the initiator's controller. All problems were run on a Linux Desktop with a 3.4GHz processor. Valid solutions were generated for all problems.

Problem	1	2	3	4
Plan Size	20	83	109	45
Solving Time (s)	4.8	12.0	34.8	35.4
Encoding Time (s)	16.2	16.1	16.3	15.9

Table 1: Plan Size, Encoding Time, and Solve Time

Problem 1: No bargaining The initiator is free to consider dialogue with the potential members, but there is one agent in particular that can achieve the goal and asks nothing in return. As expected, the final plan uses this one-agent team to achieve the goal without bargaining.

Problem 2: Birds of a feather Problem 2 is a scenario where the goal can be achieved by any of three combinations (and their supersets) of team: 1, 2-3, and 4-5. Agent 1 is capable of achieving the original goal, but if they decline, then either team 2-3 or 4-5 must be sought after. In these cases, a further restriction in the scenario is that the agents can be persuaded to join the team only if the other agent in the pair will possess a particular block in the final state (e.g., $\mathcal{B}(ag2) = \{holding\_ag3\_b1\}$ and $\mathcal{B}(ag3) = \{holding\_ag2\_b2\}$), and (e.g., $\mathcal{B}(ag4) = \{holding\_ag5\_b1\}$ and $\mathcal{B}(ag5) = \{holding\_ag4\_b2\}$). This leads the initiator to devise a plan that forms a team with one of the pairs exclusively. Note that a superset of the teams would also work, but these inefficient teams naturally are not considered by the planner: if enough of the agents have agreed to join, then a team is formed immediately. See Figure 1 for the full solution to problem 2 with the labels removed.

Problem 3: Bait and switch If the initiator acted with full honesty, only the team 1-2-5 can achieve the task in this problem. However, the initiator can find a second team that includes agent 4 instead of 5. The issue with team 1-2-4 is that agent 2 can only be persuaded if block 3 ends in room 3, while agent 4 can only be persuaded if block 3 ends in room 4. This is impossible, but the initiator's reasoning recognizes that agents 2 and 4 can both believe (one of them incorrectly) that their promise is fulfilled: in the resulting plan, agent 2 witnesses block 3 being dropped in room 3, and then the initiator directs agent 2 to walk away while agent 4 brings the block to room 4. This demonstrates the expressiveness that comes when planning with multi-agent epistemic states in lieu of the standard classical planning formalism.

Problem 4: Satisfying suspicions In problem 3, the agents continue to believe that the location of a block is unchanged even when they are in a different location. In problem 4, we change the action description so that an agent no longer believes that the location of blocks in a room remain constant when they exit (i.e., they only maintain beliefs about blocks in the room they currently inhabit). With this modification, the initiator correctly identifies just one possible team: 1-2-5. Problems 3 and 4 each take the planner approximately 35 seconds to solve, and this time largely is spent attempting to find a different configuration for dialogue acts that will result in a new team.

6 Discussion

We presented an approach for team formation from the perspective of a team initiator, whose task is to synthesize a strategy to form an effective team through capability assessment and bargaining. The team's success hinges on the members ability to achieve a collective intention that includes the original goal plus any promises made during the bargaining process.

Demonstrating and evaluating this approach on a commonly used blocksworld-style problem set, we have shown that this planning technology can handle the encoded problems readily. The relevance analysis that comes with existing planners makes our approach well suited to tackling both the dialogue and planning phases simultaneously: often the infeasibility of a team to achieve the collective intention is recognized early in the planning phase, and a new team is considered. Additionally, including the MEP formalism reveals interesting new considerations for modelling the teamwork formation problem, as evidenced by the distinction between problems 3 and 4. We intend to release our framework and test suite to the wider research community.

As the objective in this paper was not to introduce a particular mechanism for team formation, a detailed comparison with other team formation models is not relevant here. Rather, our contribution is the introduction of a novel approach that is suitable for efficiently operationalising the requirements of multiple models. Hence, we point to the generic benefits of our planning-based approach, and posit that because the underlying representation supports complex encoding, including modelling the nested beliefs of agents, many team formation models will be amenable to implementation via planning. This also includes models in other settings where collective intention is a central concept. Of course, further work, both conceptually and empirically, is needed to put a precise scope around this claim. In particular, it would be interesting to use the work of Johnson [22] to assess the performance of a team created using our approach.

The preliminary results give a good indication that automated planning techniques can solve these types of problems. Moving forward, we will expand the encoding to include richer forms of dialogue within the team formation process. In doing so, we aim to extend this work to the domain of narrative planning, where properly sequenced speech acts play a central role.

Acknowledgements: This research is partially funded by Australian Research Council Discovery Grant DP130102825, *Foundations of Human-Agent Collaboration: Situation-Relevant Information Sharing*

References

1. Agmon, N., Stone, P.: Leading ad hoc agents in joint action settings with multiple teammates. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1. pp. 341–348. IFAAMAS (2012)
2. Amgoud, L.: An argumentation-based model for reasoning about coalition structures. In: Parsons, S., Maudet, N., Moraitis, P., Rahwan, I. (eds.) Argumentation in Multi-Agent Systems, Lecture Notes in Computer Science, vol. 4049, pp. 217–228. Springer (2006)
3. Cimatti, A., Pistore, M., Roveri, M., Traverso, P.: Weak, strong, and strong cyclic planning via symbolic model checking. Artificial Intelligence 147(1), 35–84 (2003)
4. Decker, K.S., Lesser, V.R.: Designing a family of coordination algorithms. In: Proceedings of the First International Conference on Multi-Agent Systems. pp. 73–80. AAAI Press (1995)
5. Dignum, F., Dunin-Keplicz, B., Verbrugge, R.: Creating collective intention through dialogue. Logic Journal of the IGPL 9(2), 289–304 (2001)
6. Dunin-Kęplicz, B., Verbrugge, R.: Teamwork in Multi-Agent Systems: A Formal Approach. John Wiley & Sons, Ltd (2010)
7. Grant, J., Kraus, S., Perlis, D.: Formal approaches to teamwork. In: Artëmov, S.N., et al. (eds.) We Will Show Them! Essays in Honour of Dov Gabbay, Volume Two. pp. 39–68. College Publications (2005)
8. Grosz, B., Kraus, S.: Collaborative plans for complex group actions. Artificial Intelligence 86, 269–358 (1996)
9. Hoang, H., Lee-Urban, S., Muñoz-Avila, H.: Hierarchical plan representations for encoding strategic game AI. In: Proceedings of the First Artificial Intelligence and Interactive Digital Entertainment Conference, June 1-5, 2005, Marina del Rey, California, USA. pp. 63–68 (2005)
10. Hoffmann, J.: Everything you always wanted to know about planning (but were afraid to ask). In: Edelkamp, S. (ed.) Proceedings of the 34th Annual German Conference on Artificial Intelligence (KI'11). Berlin, Germany (2011)
11. Johnson, M., Jonker, C., van Riemsdijk, B., Feltovich, P., Bradshaw, J.: Joint activity testbed: Blocks world for teams (BW4T). In: Engineering Societies in the Agents World X. pp. 254–256. Springer (2009)
12. Kinny, D., Sonenberg, E., Ljungberg, M., Tidhar, G., Rao, A., Werner, E.: Planned team activity. In: Castelfranchi, C., Werner, E. (eds.) Artificial Social Systems, LNCS, vol. 830, pp. 227–256. Springer (1994)

13. Marcolino, L.S., Zhang, C., Jiang, A.X., Tambe, M.: A detailed analysis of a multi-agent diverse team. In: Balke, T., Dignum, F., van Riemsdijk, M.B., Chopra, A.K. (eds.) Coordination, Organizations, Institutions, and Norms in Agent Systems IX - COIN 2013 Revised Selected Papers. Lecture Notes in Computer Science, vol. 8386, pp. 3–24. Springer (2013)
14. Muise, C., Miller, T., Felli, P., Pearce, A., Sonenberg, L.: Efficient reasoning with consistent proper epistemic knowledge bases. In: Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (2015)
15. Muise, C.: Planner for Relevant Policies: Bitbucket repository. https://bitbucket.org/haz/pdkb-planning/ (2015), accessed: 2015-01-26
16. Muise, C., Belle, V., Felli, P., McIlraith, S., Miller, T., Pearce, A., Sonenberg, L.: Planning over multi-agent epistemic states: A classical planning approach. In: The 29th AAAI Conference on Artificial Intelligence (2015)
17. Muise, C., McIlraith, S.A., Beck, J.C.: Improved Non-deterministic Planning by Exploiting State Relevance. In: The 22nd International Conference on Automated Planning and Scheduling. The 22nd International Conference on Automated Planning and Scheduling (2012), http://www.haz.ca/papers/muise-icaps2012-fond.pdf
18. Nair, R., Tambe, M.: Hybrid BDI-POMDP Framework for Multiagent Teaming. J. Artif. Intell. Res.(JAIR) 23, 367–420 (2005)
19. Rao, A.S., Georgeff, M.P.: BDI agents: From theory to practice. In: Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95. pp. 312–319 (1995)
20. Sandholm, T., Lesser, V.: Coalitions Among Computationally Bounded Agents. Artificial Intelligence, Special Issue on Economic Principles of Multi-Agent Systems 94(1), 99–137 (1997)
21. Shehory, O., Kraus, S.: Methods for task allocation via agent coalition formation. Artificial Intelligence 101(12), 165 – 200 (1998)
22. Sierhuis, M., Jonker, C., van Riemsdijk, B., Feltovich, P.J., Bradshaw, J.M., Johnson, M.: Autonomy and interdependence in human-agent-robot teams. IEEE Intelligent Systems 27(2), 43–51 (2012)
23. Tidhar, G., Rao, A.S., Sonenberg, E.A.: Guided team selection. In: In Proceedings of the Second International Conference on Multi-Agent Systems. pp. 369–376. AAAI Press (1996)
24. Wooldridge, M., Jennings, N.R.: The cooperative problem-solving process. J. Log. Comput. 9(4), 563–592 (1999)