

Argumentation-Based Multi-agent Dialogues for Deliberation

Yuqing Tang¹ and Simon Parsons²

¹ Department of Computer Science
Graduate Center, City University of New York
365 5th Avenue, New York, NY 10016, USA
ytang@gc.cuny.edu

² Department of Computer and Information Science
Brooklyn College, City University of New York
2900 Bedford Avenue, Brooklyn, NY 11210 USA
parsons@sci.brooklyn.cuny.edu

Abstract. This paper presents an argumentation-based approach to deliberation, the process by which two or more agents reach a consensus on a course of action. The kind of deliberation that we are interested in is a process that combines both the selection of an overall goal, the reduction of this goal into sub-goals, and the formation of a plan to achieve the overall goal. We develop a mechanism for doing this, describe how this mechanism can be integrated into a system of argumentation to provide a sound and complete deliberation system, and show how the same process can be achieved through a multi-agent dialogue.

1 Introduction

Multi-agent planning is clearly an important topic for the field of multi-agents systems, and, as one might imagine, has been widely studied and for a long time. As [4] points out, there is a large variety of approaches, from distributed versions of classical AI planning techniques like NOAH [3] and partial planning [6], to techniques that were developed to exploit specific attributes of multi-agent systems like joint intentions [14,23], or the *intention-that* of SharedPlans [9]. Some of these approaches deal with multi-agent plans holistically [10], while others build plans for individual agents and then merge them [7]. Approaches as disparate as model checking [24] and auctions [26] have been adapted to generate multiagent plans.

In this paper we bring together aspects of multiagent planning and work in a field that has grown up more recently, argumentation-based dialogue [18]. While there has been much work on argumentation-based dialogue in the last few years—including that of Kraus [13], Maudet [15], McBurney [16], Reed [20], Schroeder *et al.* [21] and Sycara [22]—there is not yet a definitive account of what Walton and Krabbe [25] call *deliberation* dialogues. These are dialogues in which two or more agents converse to formulate a joint course of action.

At the time of writing, we have team formation dialogues Dignum *et al.* [5], dialogues about what should be done [8], dialogues in which one agent proposes a plan and then persuades others to adopt it [17], and even a general purpose framework for

deliberation [12]. Our goal in this paper is to provide a form of dialogue which allows agents to exchange arguments about the details of planning. This ability to debate the details is something we believe is essential if agents are going to rationally discuss what plans to adopt. In particular, we aim to develop a dialogue in which agents not only decide what to do, but create a plan *jointly*, with different sub-plans being suggested by different agents which then merge them to create an overall plan that they all agree on. This is an important step towards a complete account of deliberation.

2 Notation

As usual when considering planning, whether the classical planning of STRIPS or the decision theoretic planning of POMDPs [1], we abstract the physical world into states, actions, and state transitions caused by the actions. States and actions are the basic objects in L , the underlying language used in our approach. The kind of procedure we are interested in will determine how to compose a sequence of actions to achieve a desired state transition, namely to reach a goal from a given state.

In L , we think of a plan as being a sequence of actions, and we want to determine a plan that gets us from a specified initial state to a specified final state. The basic objects of L are:

1. A set of states: $S = \{s_0, s_1, \dots, s_n\}$.
2. A set of actions: $A = \{a_0, a_1, \dots, a_m\}$.
3. A set of pairs of states, where each pair consists of a start state and an end state. We term such a pair a *nisus*¹ and denote a set of nisi: $N = \{s_0 \hookrightarrow s'_0, s_1 \hookrightarrow s'_1, \dots, s_t \hookrightarrow s'_t\}$ where $s_i \in S$ and \hookrightarrow denotes a state transition. By way of an abbreviation, we sometimes write N as $N = \{n_0, n_1, \dots, n_t\}$ where $n_i = s_i \hookrightarrow s'_i$.
4. A set of plans: $P = \{p_0, p_1, \dots, p_s\}$. A plan p is a sequence of actions $p = a_1, \dots, a_t$ where every $a_i \in A$.

Sentences in L describe the world. In particular we are interested in what actions and plans achieve. The effect of an action can be thought of as either:

- causing a state transition; or
- achieving a nisus, denoted $a \rightsquigarrow (s \hookrightarrow s')$, where $s, s' \in S$, $a \in A$, indicating that action a achieves a transition from s to s' . This can also be written $a \rightsquigarrow n$ for nisus n .

In other words, actions bring about simple state transitions, and some of these state transitions may be distinguished as nisi — transitions between states we identify as start and end points for agents.

The effect of a plan can also be thought of in terms of state transitions: $p \rightsquigarrow (s \hookrightarrow s')$, where $p \in P$, means plan p causes a state transition from s , the *source*, to s' , the *destination*. Here p must satisfy the conditions that:

- $p = a_1, \dots, a_t$; and
- there exists a sequence of states $s_0, s_1, \dots, s_{t-1}, s_t$ such that $s = s_0$, $s' = s_t$, and $a_i \rightsquigarrow (s_{i-1} \hookrightarrow s_i)$ for $i = 1, \dots, t$.

¹ Nisus: a striving towards a goal.

Plans are thus specific sequences of actions that, when executed², will create a path through state space between two specified states. Our notion of plan is thus very much like the usual notion of a plan in simple AI planning. However, we choose to specify goals not by the usual target state but rather as a pair of initial state and final state (though we also use the conventional notion of goal in places). Why do we do this?

The answer is that the notion of a *nisi* fits rather better with our approach to deliberation than the usual notion of a goal. Obviously, plans, goals and *nisi* are all suitable for describing a state-transition graph $G = (S, A)$ where nodes S are states, edges A are assigned by atomic actions, and edges are directed³. Planning is essentially a process of finding a path between two given states, s_0 and s_g , and were we to plan in classic means-ends style, the usual notion of goal would suffice. However, we don't. The planning process we describe allows agents to work from initial state, goal state, or between any states in the middle, and in such a situation it is convenient to be able to link start and end point exactly as a *nisi* does, in order to keep track of where one is in the process.

Our slightly non-standard notation for action, $a \rightsquigarrow (s_1 \leftrightarrow s_2)$, is similarly motivated by the deliberation process. A set of action descriptions of the kind we use can be viewed as the definition of a function $a : S \rightarrow S$. By taking an action to be a function on states in this way, we capture explicitly, and in a propositional form, the fact that the same action applied to different initial states will lead to different final states (and it is a short step to capturing non-deterministic actions).

In the rest of the paper we blur the distinction between actions and plans because we take an action a to be an atomic plan — both actions and plans have the effect of creating state transitions.

3 Deliberation and Planning

We start by considering the deliberation process that a single agent goes through. The usage that Walton and Krabbe [25] make of the term “deliberation”, which is to denote the whole scope of practical reasoning, differs from that made by Bratman [2], who uses it to denote the process of choosing goals⁴ that are then subject to means-ends reasoning. Since we will be considering both types of deliberation, we will denote the first by D_{WK} , and the second by D_B . For us, the D_{WK} process starts with an overall *nisi* n_0 , and uses D_B to refine the set of sub-*nisi* in conjunction with a process of means-ends planning.

In more detail, we recursively divide D_B and the associated planning into *phases* until a plan for n_0 is reached. All the phases share the same top-level *nisi* n_0 , and each phase has a *deliberation context*. A deliberation context consists of:

1. the top-level *nisi* n_0 ,
2. a set of intermediate *nisi* N_{inter} and,
3. a set of useful plans P_{useful} (we will describe the way this set is constructed in detail below).

² Assuming, as we do for now, that actions are deterministic.

³ An action can assign more than one edge between nodes.

⁴ To be precise Bratman considers intentions not goals, but for our purposes there is little practical difference.

Within a context, an agent deliberates in the D_B sense. Based on the result of this D_B deliberation, the agent then plans. Based on the results of D_B and the subsequent planning, the agent then decides whether or not to recursively call a *child phase* to solve a sub-problem. If this is the case, the next round of D_B and planning is delayed until the child phase is complete.

To make the procedure precise, we define the following:

1. $Justified(n)$ means that a nisis n is achievable, namely it is a nisis with a plan p such that $p \rightsquigarrow n$.
2. $Src(N) = \{n | n \hookrightarrow n' \in N\}$ is the set of source states of a given set of nisis N .
3. $Dest(N) = \{s' | s \hookrightarrow s' \in N\}$ is the set of destination states of a given set of nisis N .
4. $Src(P) = \{s | p \rightsquigarrow (s \hookrightarrow s'), p \in P\}$ is the set of source states of a given set of plans P .
5. $Dest(P) = \{s' | p \rightsquigarrow (s \hookrightarrow s'), p \in P\}$ is the set of destination states of a given set of plans P .

Now, we present our first D_{WK} procedure, SD (for simple deliberation). This is called with a top-level nisis $n_0 = s_0 \hookrightarrow s_g$, initializes its top-level context with context ID $i = 0$, $N_{inter}^i = \emptyset$, and a set of partial plans that might be adopted $P_{useful}^i = A$. SD then executes the following steps:

1. Check whether $Justified(n_0)$ holds, that is whether n_0 can be achieved using plans in P_{useful}^i . If it can, then stop with a plan for n_0 .
2. Carry out D_B :
 - (a) Create a child context ID j in order to invoke a child phase.
 - (b) Choose a set of intermediate nisis N_{inter}^j for the child phase from:

Nisis of the form $s \hookrightarrow s'$ where $s \in Dest(P_{useful}^i)$ and $s' \in Src(P_{useful}^i)$.

These are all the possible state pairs which connect the end state of one existing plan to the start state of another existing plans. If there are plans for these nisis output by the planning procedure in the future, then we can create new plans by combining two existing plans with these future plans.

Nisis of the form $s_0 \hookrightarrow s'$ where $s' \in Src(P_{useful}^i)$.

These are all the possible state pairs which connect initial state of the top-level nisis to the start state of the existing plans.

Nisis of the form $s \hookrightarrow s_g$ where $s \in Dest(P_{useful}^i)$.

These are all the possible state pairs which connect the end state of the existing plans to the goal state of the top-level nisis.

Based on the candidate nisis given above, heuristics (see below) are used to gather a subset of these nisis to be N_{inter}^j , nisis that are believed to be important to achieve our top-level nisis n_0 .

Note that the last two sets of nisis are not necessary, but help to make the process more efficient.

3. Combine plans guided by the result of D_B . For each intermediate nisus $s \leftrightarrow s' \in N_{inter}^j$, combine plans as follows:
 - (a) Extend forward plans that end with initial states of nisi in N_{inter}^j : look for plans $p_1 \rightsquigarrow (s \leftrightarrow s_i)$ and $p_2 \rightsquigarrow (s_i \leftrightarrow s_j)$ and combine them to give $p_1, p_2 \rightsquigarrow (s \leftrightarrow s_j)$ if such p_1, s_i, s_j, p_2 exist.
 - (b) Extend backward plans that start with final states of nisi in N_{inter}^j : look for plans $p_1 \rightsquigarrow (s_i \leftrightarrow s_j)$ and $p_2 \rightsquigarrow (s_j \leftrightarrow s')$ and combine them to give $p_1, p_2 \rightsquigarrow (s_i \leftrightarrow s')$ if such p_1, s_i, s_j, p_2 exist.
 Add these plans into P_{useful}^i .
4. Reason about plans. Pick a subset of P_{useful}^i to be the set P_{useful}^j passed to a child phase which recursively applies the SD procedure with the new context with ID j . Repeat until all the plans in P_{useful}^i have been distributed.
5. Collect plans from child phases and combine all the P_{useful}^j into P_{useful}^i .
6. Terminate if it is clear there is no plan to achieve n_0 , otherwise go to the beginning of the procedure.

Step 2(b) is the key step in the planning process—each of the sub-parts of this step are ways in which the plan is constructed. The second and third sub-parts, respectively, capture the notions of backward chaining from the goal state and forward chaining from the initial state. The first step captures the idea that planning can work forwards and backwards simultaneously from some state in the middle of a possible plan.

Although we are describing this process for a single agent at the moment, consider how such a process might take place were several agents to be involved. In such a case different agents would be throwing out different suggestions simultaneously, and at any one time, we might have plans for achieving many different nisi “on the table”. The heuristics in the fourth sub-part of 2(b), are methods that select the most promising of such a set of nisi (which can equally well be identified by a single agent) for further consideration.

Below we will adapt this procedure first to incorporate argumentation, and then to allow it to be distributed across a pair of agents. Before we do this, we obtain soundness and completeness results:

Proposition 1 (Soundness). *If SD generates a plan p , then p is a plan to achieve the top-level nisus n_0 using the atomic actions A .*

Proof. Step 3 of the deliberation procedure ensures that only valid plans in L are composed from actions in A . Step 1 guarantees that the deliberation procedure succeeds only if there is a plan $p \rightsquigarrow n_0$. Therefore p is a valid plan to achieve n_0 using the atomic actions A . \square

Before attempting the completeness result, we need some further notation:

1. P_n is the set of plans that include n actions.
2. \oplus is a plan combination operator $P_i \oplus P_j = \{p_1, p_2 | p_1 \in P_i \wedge p_2 \in P_j\} \cup \{p_2, p_1 | p_1 \in P_i \wedge p_2 \in P_j\}$ where p_1, p_2 and p_2, p_1 must satisfy the valid plan conditions given above. This operator corresponds to step 3.

Proposition 2 (Completeness). *If there is a plan for initial nisus n_0 , then SD will succeed with a plan p which achieves n_0 .*

Proof. In step 2, P_{useful}^i determines N_{inter}^j (j is a child context of i). In step 3, N_{inter}^j determines the plans being added into P_{useful}^i . Therefore step 2 and 3 together determine the growth of P_{useful}^i . The recursive child phases called in step 5 expedite the discovery of the top-level nusus n_0 ; they don't affect the growth of P_{useful}^i . We will show that step 2 and step 3 together will grow P_{useful}^0 to contain all the plans which can be generated from atomic actions A so that if there is a plan for n_0 the deliberation procedure will certainly discover it.

We divide P_{useful}^0 into n disjoint subsets $P = P_1 \cup P_2 \cup \dots \cup P_n$, where n is the number of states, thus defining the maximum length of plans in P_n ⁵. Initially $P_1^0 = A$ and $P_i^0 = \emptyset$ for $i = 2, \dots, n$. At each point in time, steps 2 and 3 together grow P in the following way: $P_t^{k+1} = \cup(P_i^k \oplus P_j^k)$ for all $i + j = t$. Therefore if, in step 2 and 3, the P_i^k are fixed for $i = 1, \dots, t - 1$ then $P_t^{k+1} = P_t^k$. P_1 is fixed during any iteration; after the first iteration P_2 is fixed since P_1 is fixed; after the second iteration P_3 is fixed since P_1 and P_2 are fixed, and so on. In this way, after $n - 1$ iterations, P_n will be fixed. Since the maximal plan length is n , P will contain all the possible plans after $n - 1$ iterations. Therefore if there is a valid plan for n_0 , then P will contain it after $n - 1$ iterations. \square

4 Argument and Deliberation

To combine D_B with argumentation, we need to do three things. First, we extend L with predicates that control the D_B procedure. Second, we establish logic-based rules for handling nisi, reasoning about plans, combining plans and passing information through different contexts. (This will enable us to construct plans by STRIPS-like logical reasoning). Third, we add a commitment store [11] to track the course of D_B and, hence, the course of the planning process.

With a knowledge base expanded using the extended L , a plan for a nusus is certainly contained in the theorems of a subset of the knowledge base. However, the deliberation problem, to some extent, is to select an efficient way to construct a proof which backs up a plan for a nusus (the proof then becomes the justification that can be provided in a multi-agent D_{WK}). The commitment store provides a trace of how such a proof is constructed.

4.1 Additional Notation

To capture the context of phases, we introduce the following predicates into L

1. $Ultimate(n)$ denotes that $n \in N$ is the top-level nusus.
2. $N(id, n)$ denotes that $n \in N$ is an intermediate nusus in context with ID id . It is a predicate that determines whether $n \in n_{\text{inter}}^{id}$.

⁵ In any plan, we always discard any action sequence that includes a cycle $p = a_1, a_2, \dots, a_{i-1}, a_i, \dots, a_i, a_{i+1}, \dots$, because if we can construct the former plan, we will have sufficient plan fragments to eventually construct the corresponding contracted plan $p' = a_1, a_2, \dots, a_{i-1}, a_i, a_{i+1}, \dots$. Doing this effectively makes the Markov assumption, taking the effects of an action to uniquely determine the succeeding state.

3. $P(id, p)$ denotes that $p \in P$ is a useful plan in context with ID id . It is a predicate that determines whether $p \in P_{useful}^{id}$ or not.
4. $Justified(id, n)$ denotes the existence of a plan for nusus n in a the context id .
5. $Parent(id_1, id_2)$ denotes the fact that context id_1 is the parent of context id_2 .

4.2 Rules

In order to create arguments that support plans, we need to be able to trace the planning process. To do that we need to introduce the following logical rules.

Nisus Justification

$$P(i, p) \wedge [p \rightsquigarrow (s \leftrightarrow s')] \rightarrow Justified(i, s \leftrightarrow s')$$

Note that here, as in all these rules, \rightarrow denotes material implication.

Candidate Nisus Composition

$$Ultimate(n) \rightarrow N(j, n)$$

$$\begin{aligned} &Parent(i, j) \\ &\wedge P(i, p_1) \\ &\wedge P(i, p_2) \\ &\wedge [p_1 \rightsquigarrow (s' \leftrightarrow s_{foo1})] \\ &\wedge [p_2 \rightsquigarrow (s_{foo2} \leftrightarrow s)] \rightarrow N_{cand}(j, s \leftrightarrow s') \end{aligned}$$

$$\begin{aligned} &Parent(i, j) \\ &\wedge P(i, p) \\ &\wedge Ultimate(s_0 \leftrightarrow s_g) \\ &\wedge [p \rightsquigarrow (s \leftrightarrow s_{foo})] \rightarrow N_{cand}(j, s_0 \leftrightarrow s) \end{aligned}$$

$$\begin{aligned} &Parent(i, j) \\ &\wedge P(i, p) \\ &\wedge Ultimate(s_0 \leftrightarrow s_g) \\ &\wedge [p \rightsquigarrow (s_{foo} \leftrightarrow s)] \rightarrow N_{cand}(j, s \leftrightarrow s_g) \end{aligned}$$

We can use heuristics to select $N(j, n)$ from $N_{cand}(j, n)$ in order to reduce the search space. Without the heuristics, we will use the rule

$$N_{cand}(j, n) \rightarrow N(j, n)$$

so that every candidate nisus is considered.

Candidate Plan Combination

$$\begin{aligned} &[p_1 \rightsquigarrow (s \leftrightarrow s_m)] \\ &\wedge [p_2 \rightsquigarrow (s_m \leftrightarrow s')] \rightarrow p_1, p_2 \rightsquigarrow (s \leftrightarrow s') \end{aligned}$$

$$\begin{aligned}
& P(i, p_1) \\
& \wedge P(i, p_2) \\
& \wedge N(i, s \leftrightarrow s_{foo}) \\
& \wedge p_1, p_2 \rightsquigarrow (s \leftrightarrow s') \rightarrow P(i, p_1, p_2)
\end{aligned}$$

$$\begin{aligned}
& P(i, p_1) \\
& \wedge P(i, p_2) \\
& \wedge N(i, s_{foo} \leftrightarrow s) \\
& \wedge p_1, p_2 \rightsquigarrow (s \leftrightarrow s') \rightarrow P(i, p_1, p_2)
\end{aligned}$$

Plan Selection

$$Parent(i, j) \wedge P(i, p) \rightarrow P_{cand}(j, p)$$

Again we can use heuristics select $P(j, p)$ from $P_{cand}(j, p)$. Without using heuristics, we will have rule

$$P_{cand}(j, p) \rightarrow P(j, p)$$

so that every candidate plan is considered.

Plan Collection

$$P(j, p) \wedge Parent(i, j) \rightarrow P(i, p)$$

These basic rules provide a backbone to guarantee that our procedure searches the whole space of plans so that if there is a plan to achieve the nirus n_0 then we will reach it sooner or later.

4.3 Heuristics

The basic rules give us a no-frills planning procedure. Adding in heuristics like those given below tries to ensure that if there is a plan that can achieve the top-level nirus, the deliberation procedure will reach it as early as possible. We take inspiration from decision-theoretic planning [1], where choices between actions are made on the basis of their expected cost. Accordingly we introduce the following notions of cost.

1. The *action-state transition cost* $cost(a, s, s')$ is the cost of taking action a to transform state from s to s' . The value is computed or assigned outside the reasoning system.
2. The *plan-state-transition cost* $cost(p, s, s')$ is the cost of taking a plan p to transform state from s to s' . The value is computed from $cost(a, s_i, s'_{i+1})$ for all actions a in the plan p .

The *overall cost* of a plan is computed from $cost(p, s_i, s_j)$ for all the plans p that can cause state transition from s_i to s_j . We can think of the overall cost either as $cost(s, s')$, the cost of transforming state s into s' , or as $cost(i, p)$, the cost of executing plan p in context i . The idea is that although we often want to consider $p \rightsquigarrow (s \leftrightarrow s')$ as a holistic entity, to do D_B and planning we need to make comparisons between plans and actions, and we use costs to make these comparisons.

The cost of a plan can be derived from the cost of its actions in the same kind of way as it is done in decision-theoretic planning⁶. Whatever mechanism is adopted, it is outside the logical reasoning that we are studying here. Thus the assignment of costs to overall plans is, so far as the D_B and planning processes are concerned, carried out by an oracle.

Another useful notion in deciding which nisi to adopt is the *correlated valuation* of one nisi relative to another, denoted $value(n, n')$. This captures the value of achieving nisi n' in order to achieve n , and can be computed from the costs of all the plans which have the form $p \rightsquigarrow n$ for which there exists a subplan p' of p such that $p' \rightsquigarrow n'$. If there is no such sub-plan p' then $value(n, n') = 0$.

With these ideas in place, we can suggest heuristics for plan selection and nisi composition. One possibility for plan selection is to select the lowest cost plan:

$$\begin{aligned}
 &Parent(i, j) \\
 &\wedge P_{cand}(j, p) \\
 &\wedge P(i, p) \\
 &\wedge P(i, p') \\
 &\wedge cost(i, p) < cost(i, p') \rightarrow P(j, p)
 \end{aligned}$$

A possibility for nisi composition is to only adopt nisi for which the correlated valuation is above some threshold. To do this we can use:

$$\begin{aligned}
 &N(j, n) \\
 &\wedge N_{cand}(j, n') \\
 &\wedge value(n, n') > c \rightarrow N(j, n')
 \end{aligned}$$

Other heuristics can, of course, be adopted.

4.4 Single Agent Deliberation

We are now in a position to explain how a single agent can use argumentation-based D_{WK} to figure out what to do. We assume the agent has a knowledge base KB which contains a description of the physical world (e.g. the set of available actions and their effects). The agent also has a commitment store CS which it uses to trace the course of deliberation. The idea behind the procedure is to guarantee that all necessary sentences to support an argument are available in the commitment store CS before such an argument is constructed. The argumentation system used is

$$AS = \langle \mathcal{A}(KB \cup CS), Undercut, Pref \rangle$$

in the notation of [19].

The procedure for argumentation-based deliberation, SDA (Simple Deliberation through Argumentation), is given a top-level nisi $n_0 = s_0 \leftrightarrow s_g$. It first initializes the context id with $i = 0$, CS with $Ultimate(n_0)$ and $P(i, a)$ for all $a \in A$, then executes the following steps:

⁶ Such costs may be assigned by a form of reinforcement learning, for example.

1. Check \mathcal{AS} to see whether $Justified(n_0)$ is acceptable. If it is, then stop with a plan for n_0 in CS .
2. Carry out a D_B :
 - (a) Set a context ID j for a child phase.
 - (b) Using KB and CS , use \mathcal{AS} to check if $N(j, n)$ is acceptable for nisi of the following three kinds:

Nisi such as $(s \leftrightarrow s')$ for $s \in Dest(P)$ and $s' \in Src(P)$. This captures the idea of extending existing plans forwards and backwards.

Nisi such as $(s_0 \leftrightarrow s')$ for $s' \in Src(P)$. This captures the idea of extending the existing plans forwards from the source of the top-level nisis.

Nisi such as $(s \leftrightarrow s_g)$ for $s \in Dest(P)$ to capture the idea of extending the existing plans backwards from the destination of the top-level nisis.

Assert all the acceptable sentences $N(j, n)$ into CS . Notice that the above can be achieved only if the rules for nisis composition are used. If the agent exhausts all the possible candidates nisi $N_{cand}(j, n)$ but no $G(j, n)$ can be asserted, then the child phase returns to the parent context with no new plans added.

3. Combine plans guided by the results of D_B . For all the nisi $n = (s \leftrightarrow s')$ with acceptable arguments for $N(j, n)$, combine plans as follows:

Plans from s . Look for plans $p_1 \rightsquigarrow (s \leftrightarrow s_i)$ and $p_2 \rightsquigarrow (s_i \leftrightarrow s_j)$, combine them to be $p_1, p_2 \rightsquigarrow (s \leftrightarrow s_j)$ if such p_1, s_i, s_j, p_2 exist.

Plans to s' . Look for plans $p_1 \rightsquigarrow (s_i \leftrightarrow s_j)$ and $p_2 \rightsquigarrow (s_j \leftrightarrow s')$, combine them to be $p_1, p_2 \rightsquigarrow (s_i \leftrightarrow s')$ if such p_1, s_i, s_j, p_2 exist.

Check acceptability of plans with \mathcal{AS} , and assert all acceptable $p_1, p_2 \rightsquigarrow (s \leftrightarrow s_j)$, $p_1, p_2 \rightsquigarrow (s_i \leftrightarrow s')$ into CS .

4. Reason about plans. Using the plan selection rules, identify the candidate useful plans p , and for each use \mathcal{AS} to check whether $P(j, p)$ is acceptable. If it is acceptable, assert it into CS .
5. Recursively call a child phase to go through SDA with the new context with ID j .
6. Collect plans from the child phase: Use plan collection rules to identify candidates plans p . For each p , check with \mathcal{AS} whether $P(i, p)$ is acceptable. If it is acceptable, assert it into CS .
7. Go to the first step.

Since SDA is based on a process for D_B and planning that we know is sound and complete, we can easily show that it is sound and complete itself:

Proposition 3 (Soundness). *If plan p for nisis n_0 is acceptable according to \mathcal{AS} at the end of the SDA, then p achieves n_0 using actions from A .*

Proof. Step 3 combines plan p from P_1 and P_2 only if \mathcal{AS} accepts the combination. Thus \mathcal{AS} accepts the effects of p . Step 1 ensures that p is a plan for n_0 . \square

Proposition 4 (Completeness). *If there is a plan p which achieves nisis n_0 using atomic actions A , then SDA will generate p .*

Proof. Similar to the proof of Proposition 2, since steps 2 and 3 together will explore all the possible combinations of discovered acceptable plans, the procedure will generate all the plans acceptable by \mathcal{AS} at the end of the procedure. If there is plan p which achieves initial *nisi* n_0 using actions from A , then such a plan is contained in all the acceptable plans by \mathcal{AS} . \square

5 Deliberation Dialogues

We now consider how to extend the D_{WK} process to become a dialogue. We describe the dialogue process as being between just two agents, but it can easily be extended to a multi-party dialogue.

5.1 Basic Configuration

The scenario for which our deliberation dialogue was created is as follows:

1. Dialogues take place between two agents, A_1 and A_2 .
2. A_1 has initial knowledge base KB_1 and commitment store CS_1 .
3. A_2 has initial knowledge base: KB_2 and a commitment store CS_2 .
4. A_1 and A_2 share the same rules for planning and deliberation but differ in the way they evaluate plans and *nisi*. They may also have different set of actions reflecting different capabilities.
5. The context ID, id , is shared by A_1 and A_2 . Initially, $id = 0$.
6. A_1 and A_2 have an mechanism to allocate unique context IDs.
7. Both A_1 and A_2 can access CS_1 and CS_2 , hence the argumentation system of A_1 is

$$\mathcal{AS}_1 = \langle \mathcal{A}(KB_1 \cup CS_1 \cup CS_2), \text{Undercut}, \text{Pref} \rangle$$

and the argumentation system of A_2 is

$$\mathcal{AS}_2 = \langle \mathcal{A}(KB_2 \cup CS_1 \cup CS_2), \text{Undercut}, \text{Pref} \rangle.$$

Within this scenario we need to add the following idea of an auxiliary sub-dialogue.

5.2 Auxiliary Discussion Sub-dialogue

One of the reasons for agents to engage in deliberation dialogues is to combine both agents' reasoning and planning capabilities. One possible downside, the fact that conflicts may arise between the two agents, can be resolved by the use of argumentation (which, at heart, is a system for resolving conflicts in terms of the acceptability of the arguments that support the conflicting statements in two argumentation systems). To achieve this resolution we need an auxiliary discussion sub-dialogue to render a sentence acceptable to both agents, by which we mean that it is accepted by the argumentation systems \mathcal{AS}_1 and \mathcal{AS}_2 .

A discussion sub-dialogue is started by a dialogue move $discuss(p)$. Assuming that A_1 moves first, the discussion sub-dialogue proceeds as follows:

1. A_1 checks with its own argumentation system AS_1 whether p is acceptable. If it is, then A_1 makes the locution $discuss(p)$, indicating that p is open for discussion.
2. A_2 checks with its argumentation system AS_2 whether p is acceptable. If it is, then A_2 stops and declares that p is accepted by both agents. Otherwise, A_2 challenges p , indicating that it needs to see the argument for p (which will be the reason behind A_1 's suggestion of p).
3. A_1 responds to the challenge by asserting the set of support S for p .
4. For each sentence $q' \in S$, A_2 checks with AS_2 . For the unaccepted sentences $q' \in S$, A_2 discusses $\neg q'$ with A_1 , if any of the $\neg q'$ are accepted by the discussion then A_1 goes to next step; otherwise A_2 stops and declares p is accepted by both agents.
5. A_1 replaces S with another alternative support and goes back to step 3.
6. If all the possible supports for p are put forward by A_1 , but the discussion with A_2 accepts none of them, then A_1 declares that p is not accepted by both agents.

With this machinery, we can now set down the deliberation dialogue.

5.3 A Dialogue for Deliberation

Two agents can have two different set of atomic actions. This means they have different capability or different views of the physical world.

We assume that initially A_1 and A_2 agree on a top-level nius $n_0 = s_0 \leftrightarrow s_g$ (though they could arrive at this after another dialogue about what they want, a negotiation perhaps). A_1 and A_2 initialize the context id with $i = 0$. A_1 initializes CS_1 with $Ultimate(n_0)$ and $P(i, a)$ for all its actions, and A_2 does the same for its commitment store. The simple deliberation dialogue, SDD, then consists of the following steps:

1. Check a plan is required.
 A_1 discusses A_2 to check whether $Justified(n_0)$ is acceptable. If it is, then the agents stop with a plan for n_0 . Otherwise proceed to the next step.
2. Carry out D_B :
 - (a) A_1 and A_2 create a context ID j for a child phase.
 - (b) A_1 discusses with A_2 to check whether any nisi $N(j, n)$ of the following three forms are acceptable:

Nisi of the form $s \leftrightarrow s'$ where $s \in Dest(P)$ and $s' \in Src(P)$, to capture the idea of extending existing plans forwards and backwards.

Nisi of the form $s_0 \leftrightarrow s'$ where $s' \in Src(P)$, to capture the idea of extending the existing plans forwards from the source of the top-level nius.

Nisi of the form $s \leftrightarrow s_g$ where $s \in Dest(P)$, to capture the idea of extending the existing plans backwards from the destination of the top-level nius.

A_1 asserts all the acceptable sentences $N(j, n)$ into CS_1 .
 - (c) A_2 carries out step (b) but asserts results into CS_2 . Completeness hinges on doing this — otherwise some crucial information known only to A_2 might not be available (to A_1 which can use it to construct the plan).

After $A1$ and $A2$ have exhausted all the nisi n that satisfy $N_{cand}(j, n)$, and no new $N(j, n)$ are asserted, they return to the parent context.

3. Combine plans guided by the results of D_B .
 - (a) For all the nisi $n = (s \leftrightarrow s')$ with acceptable arguments for $N(j, n)$, $A1$ combines plans:

Plans from s : $A1$ looks for plans $p_1 \rightsquigarrow (s \leftrightarrow s_i)$ and $p_2 \rightsquigarrow (s_i \leftrightarrow s_j)$, and combines them to form $p_1, p_2 \rightsquigarrow (s \leftrightarrow s_j)$ if such p_1, s_i, s_j, p_2 exist.

Plans to s' : $A2$ looks for plans $p_1 \rightsquigarrow (s_i \leftrightarrow s_j)$ and $p_2 \rightsquigarrow (s_j \leftrightarrow s')$, and combines them to be $p_1, p_2 \rightsquigarrow (s_i \leftrightarrow s')$ if such p_1, s_i, s_j, p_2 exist.

$A1$ discusses $p_1, p_2 \rightsquigarrow (s \leftrightarrow s_j)$, $p_1, p_2 \rightsquigarrow (s_i \leftrightarrow s')$, with $A2$, asserting the acceptable plans into CS_1 .
 - (b) $A2$ does the same as $A1$ does for (a) but asserts the results into CS_2 .
4. Reason about plans:
 - (a) $A1$ uses the plan selection rules to identify candidate useful plans p , and for such plans discusses with $A2$ whether $P(j, p)$ is acceptable. If one or more p are acceptable, $A1$ asserts them into CS_1 .
 - (b) $A2$ carries out the analogous process.
5. Recursively call child phases to go through SDD with a new context.
6. Collect plans from the child phases:
 - (a) $A1$ uses the plan collection rules to figure out which candidate plans p should be collected. Then it discusses with $A2$ to check whether the resulting $P(i, p)$ are acceptable. If one or more are acceptable, then they are asserted into CS_1 .
 - (b) $A2$ carries out the analogous process
7. Go to the beginning of the procedure.

Once again we can prove the soundness and completeness of the procedure, showing that recasting it as a dialogue does not detract from it:

Proposition 5 (Soundness). *If plan p for nusus n_0 is acceptable by both agents at the end of SDD, then p achieves n_0 using atomic actions that both agents agree upon.*

Proof. Step 3 combines plan p from P_1 and P_2 only if AS_1 and AS_2 both accept the combination. Thus both agents agree on the effects of p . Step 1 ensures that both agents agree that p is a plan for n_0 . \square

Proposition 6 (Completeness). *If there is a plan p which achieves nusus n_0 using a set of atomic actions that both agents agree upon, then SDD will generate p .*

Proof. Similar to the proof of Propositions 2 and 4, since steps 2 and 3 together will explore all the possible combinations of discovered plans accepted by both agents, the procedure will generate all the plans acceptable by both AS_1 and AS_2 at the end of the procedure. If, according to the acceptable atomic actions A agreed by both agents, there is plan p which achieves initial nusus n_0 , then such a plan is contained in all the acceptable plans by AS_1 and AS_2 . \square

6 Discussion

The full argumentation-based deliberation dialogue SDD successfully composes plans in one of the following ways:

1. *A1* composes the whole plan, *A2* agrees with it.
2. *A2* composes the whole plan, *A1* agrees with it.
3. *A1* composes some parts of the plan; *A2* composes some other parts of the plan; *A1* and *A2* combine the two parts to create a partial plan that both *A1* and *A2* agree on, and so on until the whole plan is constructed.

Thus we can see that our process combines the “create a plan and then convince others it works” approach of [17] with the “merge different plans” approach of [7]. For a given situation, SDD will typically take an approach that is a mixture of the two, involving the creation and merging of separate sub-plans, some of which may be as simple as a single action. Some of this merging will involve one agent persuading the other to adopt the sub-plan. SDD thus achieves the overall goal that we set out at the start of the paper.

One thing to note about this work concerns the heuristics used to guide the search for plans during deliberation. We never specified these in any detail, though we give some high-level hints about the possible form that they may take. This does not detract from the formal results, since the results hold even if we have no heuristics (in which case we essentially do an exhaustive search through the full state-space). However, decent heuristics will help to focus the search and thus make it more efficient.

Finally, we should note that there are clearly some similarities between the way in which our approach to deliberation allows plans to grow around any *nisus* (that is to grow forwards, backwards, or in both directions from somewhere in what turns out to be the middle of the plan) and partial order planning. This relationship is something we will explore in the future. It is particularly thought-provoking since the part of this work that we believe is most valuable, that is the integration of planning and multi-agent argumentation-based dialogue, does not depend upon the precise kind of planning used (we simply picked one form of planning to make our suggestion concrete). As a result, partial order planning could easily be fitted into our framework, and it would be interesting to investigate the kind of system that resulted from doing this.

7 Conclusion

This paper has described a mechanism for carrying out deliberation dialogues in the sense of Walton and Krabbe [25] at a high level of detail — that is dialogues in which agents decide what actions to perform and the order in which they are performed. Our approach, which as described is limited to two agents but could easily be generalized, recursively mixes *nisus* selection and planning, allowing these tasks to be distributed between the agents in a flexible way. The approach makes it possible for agents to combine their knowledge about the environment, and to make use of the planning abilities of both agents (since one can readily imagine that they have complementary expertise, as embodied in the heuristics they can employ).

Two directions of future research are particularly attractive to us. First, as mentioned above, it seems appropriate to allow the agents to learn the values of actions across a

number of trials, and this might easily be achieved by techniques from reinforcement learning. Doing so suggests a bridge between the kind of procedure we have developed here and multi-agent decision theoretic planning of the kind considered in [10]. Exploring such connections is the second direction we intend to take.

Acknowledgements. This work was made possible by funding from NSF #REC-02-19347, NSF #IIS 0329037 and EU FP6-IST 002307 (ASPIC).

References

1. Craig Boutilier, Thomas Dean, and Steve Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.
2. M. E. Bratman. *Intention, plans, and practical reason*. CSLI Publications, 1999.
3. D. D. Corkill. Hierarchical planning in a distributed environment. In *Proceedings of the International Joint Conference on Artificial Intelligence*, pages 168–175, 1979.
4. M. E. desJardins, E. H. Durfee, C. L. Ortiz, and M. J. Wolverton. A survey of research in distributed, continual planning. *AI Magazine*, 21(1), 2000.
5. F. Dignum, B. Dunin-Kępicz, and R. Verbrugge. Agent theory for team formation by dialogue. In C. Castelfranchi and Y. Lespérance, editors, *Seventh Workshop on Agent Theories, Architectures, and Languages*, pages 141–156, Boston, USA, 2000.
6. E. H. Durfee and V. R. Lesser. Partial global planning: A coordination framework for distributed hypothesis formation. *IEEE Transactions on Systems, Man, and Cybernetics*, KDE-1:63–83, 1991.
7. E. Ephrati and J. S. Rosenschein. Multi-agent planning as the process of merging distributed sub-plans. In *Proceedings of the Twelfth International Workshop on Distributed Artificial Intelligence*, pages 115–129, 1993.
8. K. Greenwood, T. Bench-Capon, and P. McBurney. Structuring dialogue between the People and their representatives. In R. Traumnüller, editor, *Electronic Government: Proceedings of the Second International Conference (EGOV03), Prague, Czech Republic*, Lecture Notes in Computer Science 2739, pages 55–62, Berlin, Germany, 2003. Springer.
9. B. J. Grosz and S. Kraus. The evolution of SharedPlans. In M. J. Wooldridge and A. Rao, editors, *Foundations of Rational Agency*, volume 14 of *Applied Logic*. Kluwer, The Netherlands, 1999.
10. C. Guestrin, D. Koller, and R. Parr. Multiagent planning with factored MDPs. In *Advances in Neural Information Processing Systems*, pages 1523–1530, 2001.
11. C. L. Hamblin. *Fallacies*. Methuen and Co Ltd, London, UK, 1970.
12. D. Hitchcock, P. McBurney, and S. Parsons. The eightfold way of deliberation dialogues. *International Journal of Intelligent Systems*, 2004.
13. S. Kraus, K. Sycara, and A. Evenchik. Reaching agreements through argumentation: a logical model and implementation. *Artificial Intelligence*, 104(1–2):1–69, 1998.
14. H. Levesque, P. Cohen, and J. Nunes. On acting together. In *Proceedings of the National Conference on Artificial Intelligence*, pages 94–99, 1990.
15. N. Maudet and F. Evrard. A generic framework for dialogue game implementation. In *Proceedings of the 2nd Workshop on Formal Semantics and Pragmatics of Dialogue*, University of Twente, The Netherlands, May 1998.
16. P. McBurney. *Rational Interaction*. PhD thesis, University of Liverpool, 2002.
17. P. Panzarasa, N. R. Jennings, and T. J. Norman. Formalising collaborative decision making and practical reasoning in multi-agent systems. *Journal of Logic and Computation*, 12(1):55–117, 2002.

18. S. Parsons and P. McBurney. Argumentation-based dialogues for agent coordination. *Group Decision and Negotiation*, 12(5):415–439, 2003.
19. S. Parsons, M. Wooldridge, and L. Amgoud. An analysis of formal inter-agent dialogues. In *1st International Conference on Autonomous Agents and Multi-Agent Systems*. ACM Press, 2002.
20. C. Reed. Dialogue frames in agent communications. In Y. Demazeau, editor, *Proceedings of the Third International Conference on Multi-Agent Systems*, pages 246–253. IEEE Press, 1998.
21. M. Schroeder, D. A. Plewe, and A. Raab. Ultima ratio: should Hamlet kill Claudius. In *Proceedings of the 2nd International Conference on Autonomous Agents*, pages 467–468, 1998.
22. K. Sycara. Argumentation: Planning other agents' plans. In *Proceedings of the Eleventh Joint Conference on Artificial Intelligence*, pages 517–523, 1989.
23. M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83–124, 1997.
24. W. van der Hoek and M. Wooldridge. Tractable multiagent planning for epistemic goals. In *Proceedings of the 1st International Conference on Autonomous Agents and Multiagent Systems*, 2002.
25. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany, NY, USA, 1995.
26. R. Zlot and A. Stentz. Market-based multirobot coordination using task abstraction. In *Proceedings of the 4th International Conference on Field and Service Robotics*, 2003.