Deliberation Dialogues During Multi-Agent Planning

Barbara Dunin-Kęplicz[1], Alina Strachocka[2] and Rineke Verbrugge[3]

[1] Institute of Informatics, Warsaw University, Warsaw, Poland,
and ICS, Polish Academy of Sciences, Warsaw, Poland
keplicz@mimuw.edu.pl,
[2] Institute of Informatics, Warsaw University, Warsaw, Poland
astrachocka@mimuw.edu.pl
[3] Department of Artificial Intelligence, University of Groningen,
Groningen, The Netherlands, rineke@ai.rug.nl

Abstract. Cooperation in multi-agent systems essentially hinges on appropriate communication. This paper shows how to model communication in teamwork within TeamLog, the first multi-modal framework wholly capturing a methodology for working together. Starting from the dialogue theory of Walton and Krabbe, the paper focuses on deliberation, the main type of dialogue during team planning. We provide a schema of deliberation dialogue along with semantics of adequate speech acts, this way filling the gap in logical modeling of communication during planning.

1 Introduction

Typically teamwork in multi-agent systems (MAS) is studied in the context of BGI (*Beliefs, Goals* and *Intentions*, commonly called BDI) systems, allowing extensive reasoning about agents' informational and motivational attitudes necessary to work together. Along this line, TeamLog [5], a framework for modeling teamwork, has been created on the basis of multi-modal logic. It provides rules for establishing and maintaining a cooperative team of agents, tightly bound by a collective intention and working together on the basis of collective commitment.

Although communication schemes during teamwork were formulated as an inherent part of TeamLog [4], this aspect of multi-agent planning was not yet treated in detail. To fill the gap, a model of deliberation dialogue during planning is investigated in this research. When a team collectively intends to achieve a goal, it needs to decide how to divide this into subgoals, to choose a sequence of actions realizing them, and finally to allocate the actions to team members. We structure these phases as deliberation dialogues, accompanied by ongoing belief revision. Thus, we formally model the team's important transition from a collective intention, to a plan-based social commitment, making it ready for action.

The paper is organized as follows. Sections 2 and 3 briefly introduce speech acts, dialogue, and teamwork theory. Next, in Section 4 the logical language is given, followed by discussion of the consequences of speech acts in Section 5. Sections 6 and 7, the heart of the paper, introduce a new model of deliberation and elaborate on planning. Finally, conclusions and plans for future work are presented.

2 Speech Acts and Dialogues

Communication in MAS has two pillars: Walton and Krabbe's semi-formal theory of dialogue [16] and the speech acts theory of Austin and Searle [15, 3]. Walton and Krabbe identified six elementary types of dialogues: *persuasion*, *negotiation*, *inquiry*, *information seeking*, *eristics* and, central to this paper, *deliberation*.

Deliberation starts from an open, practical problem: a need for action. It is often viewed as agents' *collective* practical reasoning, where they determine which goals to attend and which actions to perform. While dialogues can be seen as the building blocks of communication, they in turn are constructed from *speech acts*.

Research on speech acts belongs to philosophy of language and linguistics since the early 20th century. The basic observation of Austin [3], that some utterances cannot be verified as true or false, led to the division of speech acts into *constatives*, which can be assigned a logical truth value, and the remaining group of *performatives*. The second father of speech acts theory, Searle, created their most popular taxonomy, identifying: *assertives*, committing to the truth of a proposition (e.g., stating), *directives*, which get the hearer to do something (e.g., asking), *commissives*, committing the speaker to some future action (e.g., promising), *expressives*, expressing a psychological state (e.g., thanking), and *declaratives*, which change reality according to the proposition (e.g., baptising).

Speech acts theory has been extensively used in modeling communication in MAS to express intentions of the sender [8]. There have been many approaches to defining their semantics [13, 2, 12, 9], still some researchers view them as primitive notions [14]. Within the most popular *mentalistic* approach, reflected in languages such as KQML and FIPA ACL [8], speech acts are defined through their impact on agents' mental attitudes. The current paper clearly falls therein (see especially Section 5). Let us place dialogues in the context of teamwork.

3 Stages of Teamwork

In multi-agent cooperative scenarios, communication is inevitable and teamwork, as the pinnacle of cooperation, plays a vital role. The common division of teamwork into four stages originates from [17], while a complete model, binding these stages to formalized team attitudes, can be found in [5]. In summary:

1. Potential recognition. Teamwork begins when an initiator needs assistance and looks for potential groups of agents willing to cooperate to achieve a goal.

2. During **team formation** a loosely-coupled group of agents is transformed into a strictly cooperative team sharing a *collective intention* towards the goal (φ).

3. During **plan formation** a team *deliberates* together how to proceed, concluding in a collective commitment, based on a *social plan*. Collective planning consists of the three phases: *task division*, leading to $division(\varphi, \sigma)$ (see table 1); *means-end analysis*, leading to $means(\sigma, \tau)$, and *action allocation*, leading to $allocation(\tau, P)$. Success of these phases is summed up by $constitute(\varphi, P)$.

4. During **team action** agents execute their share of the plan. In real situations, many actions are at risk of failure, calling for a necessary reconfiguration [5], that amounts to the intelligent and situation-sensitive replanning.

With each stage of teamwork, adequate notions in TeamLog are connected. As there is no room for discussing them in detail, please see [5].

$\mathrm{BEL}(i,\varphi)$	agent i believes that φ
$\text{E-BEL}_G(\varphi)$	all agents in group G believe φ
$\text{C-BEL}_G(\varphi)$	group G has the common belief that φ
$\mathrm{GOAL}(a,\varphi)$	agent a has the goal to achieve φ
$\mathrm{INT}(a,\varphi)$	agent a has the intention to achieve φ
$\mathrm{COMM}(i,j,\alpha)$	agent i commits to j to perform α
$\textit{do-ac}(i,\alpha)$	agent i is just about to perform action α
$\textit{division}(\varphi,\sigma)$	σ is the sequence of subgoals resulting from decomposition of φ
$\textit{means}(\sigma,\tau)$	τ is the sequence of actions resulting from means-end analysis on σ
$\textit{allocation}(\tau,P)$	P is a social plan resulting from allocating the actions from τ
$\textit{constitute}(\varphi,P)$	P is a correct social plan for achieving φ
$\texttt{confirm}(\varphi)$	plan to test if φ holds at the given world
$\mathrm{prefer}(i,x,y)$	agent i prefers x to y

Table 1. Formulas and their intended meaning

4 The logical language

We introduce a subsystem of TeamLog$^{\mathrm{dyn}}$ (see [5, Chapters 5 and 6]), containing solely the elements crucial to team planning. Individual actions and formulas are defined inductively.

Definition 1. The language is based on the following sets:
- a denumerable set $\mathcal{P}$ of *propositional symbols*;
- a finite set $\mathcal{A}$ of *agents*, denoted by $1, 2, \ldots, n$;
- a finite set $\mathcal{A}t$ of *atomic actions*, denoted by a or b.

In TeamLog most modalities expressing agents' motivational attitudes appear in two forms: with respect to *propositions* reflecting a particular state of affairs, or with respect to *actions*. The set of formulas $\mathcal{L}$ (see Definition 4) is defined by a simultaneous induction, together with the set of individual actions $\mathcal{A}c$ and the set of social plan expressions $\mathcal{S}p$ (see Definitions 2 and 3). Individual actions may be combined into group actions by the social plan expressions.

Definition 2. The set $\mathcal{A}c$ of individual actions is defined inductively as follows:
AC1 each atomic action $a \in \mathcal{A}t$ is an individual action;
AC2 if $\varphi \in \mathcal{L}$, then $\texttt{confirm}(\varphi)$ is an individual action[4];
AC3 if $\alpha_1, \alpha_2 \in \mathcal{A}c$, then $\alpha_1; \alpha_2$ is an individual action, standing for α_1 followed by α_2;
AC4 if $\alpha_1, \alpha_2 \in \mathcal{A}c$, then $\alpha_1 \cup \alpha_2$ is an individual action, standing for nondeterministic choice between α_1 and α_2;
AC5 if $\alpha \in \mathcal{A}c$, then α^* is an individual action, standing for "repeat α a finite, but nondeterministically determined, number of times";
AC6 if $\varphi \in \mathcal{L}$, $i, j \in \mathcal{A}$ and $G \subseteq \mathcal{A}$, then the following are individual actions: $\texttt{announce}_{i,G}(\varphi)$, $\texttt{assert}_{i,j}(\varphi)$, $\texttt{request}_{i,j}(\varphi)$, $\texttt{concede}_{i,j}(\varphi)$.

[4] In PDL, $\texttt{confirm}(\varphi)$ is usually denoted as "φ?", standing for "proceed if φ is true, else fail".

In addition to the standard dynamic operators of [**AC1**] to [**AC5**], the communicative actions of [**AC6**] are introduced. For their meanings, see Section 5. Their interplay is the main matter of this paper.

Definition 3. The set $\mathcal{S}p$ of social plan expressions is defined inductively:
SP1 If $\alpha \in \mathcal{A}c$, $i \in \mathcal{A}$, then $do\text{-}ac(i, \alpha)$ is a well-formed social plan expression;
SP2 If $\varphi \in \mathcal{L}$, then $\texttt{confirm}(\varphi)$ is a social plan expression;
SP3 If α and β are social plan expressions, then $\langle\alpha;\beta\rangle$ (sequential composition) and $\langle\alpha \parallel \beta\rangle$ (parallellism) are social plan expressions.

Definition 4. The set of formulas $\mathcal{L}$ is defined inductively:
F1 each atomic proposition $p \in \mathcal{P}$ is a formula;
F2 if $\varphi, \psi \in \mathcal{L}$, then so are $\neg\varphi$ and $\varphi \wedge \psi$;
F3 if $\varphi \in \mathcal{L}$, $\alpha \in \mathcal{A}c$, $i, j \in \mathcal{A}$, $G \subseteq \mathcal{A}$, $\sigma, \sigma_1, \sigma_2$ are finite sequences of formulas, τ is a finite sequence of individual actions, and $P \in \mathcal{S}p$ is a social plan expression, then the following are formulas:
epistemic modalities $\mathrm{BEL}(i, \varphi)$, $\text{E-BEL}_G(\varphi)$, $\text{C-BEL}_G(\varphi)$;
motivational modalities $\mathrm{GOAL}(i, \varphi)$, $\mathrm{INT}(i, \varphi)$, $\text{E-INT}_G(\varphi)$, $\text{M-INT}_G(\varphi)$, $\text{C-INT}_G(\varphi)$, $\mathrm{COMM}(i, j, \alpha)$, $\text{S-COMM}_{G,P}(\varphi)$;
execution modalities $do\text{-}ac(i, \alpha)$;
stage results $division(\varphi, \sigma)$, $means(\sigma, \tau)$, $allocation(\tau, P)$, $constitute(\varphi, P)$;
other $\mathrm{PROOF}(\varphi)$, $\mathrm{prefer}(i, \sigma_1, \sigma_2)$.
Epistemic and motivational modalities are governed by the axioms given in the Appendix.

The predicate $constitute(\varphi, P)$ stands for "P is a correctly constructed social plan to achieve φ". Formally:

$$constitute(\varphi, P) \leftrightarrow \bigvee_{\sigma} \bigvee_{\tau} (division(\varphi, \sigma) \wedge means(\sigma, \tau) \wedge allocation(\tau, P))$$

A team G has a *mutual intention* to achieve goal φ ($\text{M-INT}_G(\varphi)$) if all intend it, all intend that all intend it, and so on, ad infinitum. To create a *collective intention* ($\text{C-INT}_G(\varphi)$), a common belief about the mutual intention should be established during team formation. Then, during plan formation, the team chooses a social plan P to achieve φ. On its basis, they create a *collective commitment* ($\text{S-COMM}_{G,P}(\varphi)$), including team members' social commitments ($\mathrm{COMM}(i, j, \alpha)$) to perform their allocated actions. The axiom system providing definitions for these notions can be found in the Appendix.

4.1 Kripke Models

Each Kripke model for the language defined above consists of a set of worlds, a set of accessibility relations between worlds, and a valuation of the propositional atoms. The definition also includes semantics for derived operators corresponding to performance of individual actions.

Definition 5. A Kripke model is a tuple:

$$\mathcal{M} = (W, \{B_i : i \in \mathcal{A}\}, \{G_i : i \in \mathcal{A}\}, \{I_i : i \in \mathcal{A}\}, \{R_{i,\alpha} : i \in \mathcal{A}, \alpha \in \mathcal{A}c\}, Val, nextac),$$

such that

1. W is a set of possible worlds, or states;
2. For all $i \in \mathcal{A}$, it holds that $B_i, G_i, I_i \subseteq W \times W$. They stand for the accessibility relations for each agent w.r.t. beliefs, goals, and intentions, respectively.
3. For all $i \in \mathcal{A}$, $\alpha \in \mathcal{A}c$, it holds that $R_{i,\alpha} \subseteq W \times W$. They stand for the dynamic accessibility relations. Here, $(w_1, w_2) \in R_{i,\alpha}$ means that w_2 is a possible resulting state from w_1 by i executing action α.
4. $Val : \mathcal{P} \times W \to \{0,1\}$ is the function that assigns the truth values to propositional formulas in states.
5. $nextac : \mathcal{A} \times \mathcal{A}c \to (W \to \{0,1\})$ is the next moment individual action function such that $nextac(i,\alpha)(w)$ indicates that in world w agent i will next perform action α. $\mathcal{M}, v \models do\text{-}ac(i,\alpha) \Leftrightarrow nextac(i,\alpha)(v) = 1$.

In the semantics, the relations $R_{i,a}$ for atomic actions a are given. The other accessibility relations $R_{i,\alpha}$ for actions are built up from these in the usual way [10]:

$$(v,w) \in R_{i,\texttt{confirm}(\varphi)} \Leftrightarrow (v = w \text{ and } \mathcal{M}, v \models \varphi);$$
$$(v,w) \in R_{i,\alpha_1;\alpha_2} \Leftrightarrow \exists u \in W[(v,u) \in R_{i,\alpha_1} \text{ and } (u,w) \in R_{i,\alpha_2}];$$
$$(v,w) \in R_{i,\alpha_1 \cup \alpha_2} \Leftrightarrow [(v,w) \in R_{i,\alpha_1} \text{ or } (v,w) \in R_{i,\alpha_2}];$$
R_{i,α^*} is the reflexive and transitive closure of $R_{i,\alpha}$.

Definition 6. Let $\varphi \in \mathcal{L}$, $i \in \mathcal{A}$ and $\alpha \in \mathcal{A}c$.

$$\mathcal{M}, v \models [do(i,\alpha)]\varphi \Leftrightarrow \text{ for all } w \text{ with } (v,w) \in R_{i,\alpha}, \mathcal{M}, w \models \varphi.$$

For the dynamic logic of actions, we adapt the axiomatization of propositional dynamic logic (PDL) [10]. The system described above has an EXPTIME-hard decision problem, just like TeamLog$^{\text{dyn}}$ and TeamLog itself [7].

5 Semantics of Speech Acts

In TeamLog, deliberation is modeled via elementary speech acts `assert`, `concede` and `request`, and the compound speech acts `challenge` and `announce`, defined in terms of PDL and described before in [4]. They are treated as ordinary actions and distinguished by their consequences. Utterances often necessitate participants' belief revision, which may be handled by diverse methods (see [1]).

In the sequel, the construction "**if** φ **then** α **else** β" abbreviates the PDL expression $(\texttt{confirm}(\varphi);\alpha) \cup (\texttt{confirm}(\neg\varphi);\ \beta)$, and analogously for "**if** φ **then** α", where $\texttt{confirm}(\varphi)$ refers to testing whether φ holds (see [5, Chapter 6]). The construct $[\beta]\varphi$ means that after performing β, φ holds.

Consequences of assertions $\texttt{assert}_{a,i}(\varphi)$ stands for agent a telling agent i that φ holds. According to the fundamental assumption that agents are as truthful as they can be, each $\texttt{assert}(\varphi)$ obliges the sender to believe in φ.

Definition 7. The consequences of assertions:

CA $[\texttt{assert}_{a,i}(\varphi)]\ (\text{BEL}(i,\varphi) \wedge\ \text{BEL}(i, \text{BEL}(a,\varphi)))$

The recipient has two possibilities to react. Unless having beliefs conflicting with φ, it answers with a $\texttt{concede}_{i,a}$. Otherwise, with a $\texttt{challenge}_{i,a}$:

$$\neg\text{BEL}(i,\neg\varphi) \to do\text{-}ac(i, \texttt{concede}_{i,a}(\varphi))$$
$$\text{BEL}(i,\neg\varphi) \to do\text{-}ac(i, \texttt{challenge}_{i,a}(\varphi))$$

Consequences of requests $\texttt{request}_{a,i}(\alpha)$ stands for agent a requesting agent i to perform the action α. The sender, after requesting information about φ (with $\alpha = \texttt{assert}_{i,a}(\varphi)$), must wait for a reply. The receiver i has four options:

1. To ignore a and not answer at all.
2. To state that it is not willing to divulge this information.
3. To state that it does not have enough information about φ:
$$\texttt{assert}_{i,a}(\neg(BEL(i,\varphi) \wedge \neg BEL(i,\neg\varphi))).$$
4. Either to assert that φ is the case or that it is not:
$\mathrm{BEL}(i,\varphi) \to do\text{-}ac(i, \texttt{assert}_{i,a}(\varphi))$ and $\mathrm{BEL}(i,\neg\varphi) \to do\text{-}ac(i, \texttt{assert}_{i,a}(\neg\varphi))$.

The consequences are the same as for proper assertions.

Consequences of concessions $\texttt{concede}_{a,i}(\varphi)$ stands for agent a's communicating its positive attitude towards φ to i. Concessions are similar to assertions. The only difference is that i can assume that a believes φ in the course of dialogue, but might retract it afterwards.

Definition 8. The consequences of concessions:
CCO $[\texttt{concede}_{a,i}(\varphi)]\mathrm{BEL}(i, \mathrm{BEL}(a,\varphi))$.

Consequences of challenges $\texttt{challenge}_{a,i}(\varphi)$ stands for a's communicating its negative attitude towards φ to i. The consequences of `challenge` are more complicated due to the complexity of the speech act itself. It consists of a negation of φ and of a request to prove φ.

Definition 9. If $\varphi, PROOF(\varphi) \in \mathcal{L}$, $a, i \in \mathcal{A}$, then
CH $\texttt{challenge}_{a,i}(\varphi) \equiv \texttt{assert}_{a,i}(\neg\varphi); \texttt{request}_{a,i}(\texttt{assert}_{i,a}(PROOF(\varphi)))$

The answer to the `request` in `challenge` should comply with the rules above. If i can prove φ, it should answer with speech act $\texttt{assert}_{i,a}(PROOF(\varphi))$ being committed to $PROOF(\varphi)$. In return, a should refer to i's previous answer. Thus[5], the consequences of `challenge` depend on the dialogue and can be twofold.

Definition 10. The consequences of challenges:
CH1 $[\texttt{challenge}_{a,i}(\varphi)]\ (\mathrm{BEL}(a,\varphi) \wedge \mathrm{BEL}(i, \mathrm{BEL}(a,\varphi)) \wedge \mathrm{BEL}(a, PROOF(\varphi)) \wedge \mathrm{BEL}(i, \mathrm{BEL}(a, PROOF(\varphi))))$
CH2 $[\texttt{challenge}_{a,i}(\varphi)]\ (\neg\mathrm{BEL}(i,\varphi) \wedge \mathrm{BEL}(a, \neg\mathrm{BEL}(i,\varphi)) \wedge \neg\mathrm{BEL}(i, PROOF(\varphi)) \wedge \mathrm{BEL}(a, \neg\mathrm{BEL}(i, PROOF(\varphi))))$

In first case, **[CH1]**, a admits it was wrong. The agents' beliefs have changed, reflected by the acceptance of i's proof, which led to belief revision about φ. In the second case, i admits it was wrong. Belief revision regarding rejecting the proof of φ leads to updating beliefs about φ.

Consequences of announcements An announcement $\texttt{announce}_{a,G}(\varphi)$ can be seen as a complex assertion standing for "agent a announces to group G that φ holds". In addition, the agent passes a message that the same information has been delivered to the whole group. The group becomes commonly aware that φ.

Definition 11. Consequences of announcements:
CAN $[\texttt{announce}_{a,G}(\varphi)]\ \text{C-BEL}_G(\varphi)$.

Once the logical language is set, we may proceed to the core of this paper.

[5] Assuming the rule $\mathrm{BEL}(a, PROOF(\varphi)) \to \mathrm{BEL}(a,\varphi)$

6 A Four-stage Model of Deliberation

The schema for deliberation dialogues presented below benefits from the model of McBurney, Hitchcock and Parsons [11]. It starts from a formal opening, introducing the subject of the dialogue, aiming to make a common decision, confirmed in a formal closure. Deliberation on "$\psi(x)$"[6] aims at finding the best t satisfying ψ from a finite candidate set T_ψ and to create a common belief about this among the team. Even though deliberation during teamwork is a collective activity, its structure is imposed by the initiator a. Other agents follow the rules presented below. Failure at any of the dialogue stages causes backtracking.

Opening Agent a's first step is to open the deliberation dialogue on the subject ψ by a request to all other $i \in G$:

$$\texttt{request}_{a,i}(\textbf{ if } \bigvee_{t \in T_\psi} \psi(t) \textbf{ then } \texttt{assert}_{i,a}(\psi(t))\textbf{else } \texttt{assert}_{i,a}(\neg \bigvee_{t \in T_\psi} \psi(t)))$$

As always after requests, agents have four ways of answering (see Section 5). If no one answers, deliberation fails. Agent a waits for a certain amount of time before concluding on the answers from group G.

Voting During voting, a announces to all $i \in G$ its finite set $T_{\psi,a}$ of all or preselected answers collected before:

$$\texttt{assert}_{a,i}(\bigwedge_{t \in T_{\psi,a}} \bigvee_{i \in G} \mathrm{BEL}(i, \psi(t)))$$

In words, a asserts to every agent that for each preselected answer t, there is an agent i in the group believing that $\psi(t)$ is the case.
Next, agent a opens the voting by a request to all $i \in G$:

$$\texttt{request}_{a,i}(\bigwedge_{x,y \in T_{\psi,a}} (\textbf{ if } \mathrm{BEL}(i, \psi(x)) \wedge \mathrm{prefer}(i, x, y) \textbf{ then } \texttt{assert}_{i,a}(\mathrm{prefer}(i, x, y))))$$

If no one answers, the scenario leads back to step 1, which is justified because the communication in step 2 may entail some belief revisions. Should some answers be received, a "counts the votes", possibly using different evaluation functions.

Confirming Then, a announces the winning proposal w and requests all opponents from G to start a persuasion:

$$\begin{aligned}\texttt{request}_{a,i}(&\textbf{if } \mathrm{BEL}(i, \neg\psi(w)) \vee \textstyle\bigvee_{t \in T_{\psi,a}} (\mathrm{prefer}(i, t, w)) \\ &\textbf{then } \texttt{assert}_{i,a}(\neg\psi(w) \vee \mathrm{prefer}(i, t, w)))\end{aligned}$$

During this phase, if no agent steps out, the scenario moves to the closure. If, on the other hand, there is an agent j who thinks that w is not the best option, it has to announce this and challenge a to provide a proof (using `challenge`). Thus the dialogue switches to persuasion, where j must convince a of the competing offer t, or that $\psi(w)$ doesn't hold. If it succeeds, a adopts and heralds agent's j thesis to all $i \in G$:

$$\texttt{assert}_{a,i}(\neg\psi(w) \vee \mathrm{prefer}(a, t, w))$$

[6] ψ is usually ungrounded, e.g. $\psi(x) = \mathit{president}(x)$. The answers are (partially) grounded terms, e.g., $\mathit{president}(\mathit{JohnSmith})$.

In this situation, the remaining agents may concede or may challenge the thesis:

$$\texttt{concede}_{i,a}(\neg\psi(w) \vee \text{prefer}(a,t,w)) \text{ or } \texttt{challenge}_{i,a}(\neg\psi(w) \vee \text{prefer}(a,t,w)).$$

If they choose to challenge, a must get involved into persuasion with the challenging agent. Finally, when all conflicts have been resolved, the scenario moves on.

Closure At last, a announces the final decision w: $\texttt{announce}_{a,G}(\psi(w))$.

Deliberating agents collaborate on the future course of actions, each of them trying to influence the final outcome. The principal kind of reasoning here is goal-directed practical reasoning, leading to a plan.

7 Unveiling the Plan Formation Stage

The (formal) aim of plan formation is transition from a collective intention to a collective commitment (see Section 3), achieved by means of dialogue. Consider, as an example, a team of various robots: digger (D), truck (T) and team leader (L) with a goal to restore order after a building collapses: $\varphi = \texttt{order}$. They discover another working team soon: first aid (FA) and 20 swarm robots ($S_1, \ldots, S_{20}$) and decide to join forces to better serve their goal. Suppose potential recognition and team formation have been successful. Then, the first phase of planning, *task division*, aims at dividing the overall goal φ into a new sequence of subgoals. L opens the deliberation dialogue by requesting all other $i \in G$ to share their ideas:

$$\begin{aligned}\texttt{request}_{L,i}(\ \textbf{if}\ \textstyle\bigvee_{\sigma\in T_{Goals}} \mathit{division}(\texttt{order},\sigma)\ \textbf{then}\ \texttt{assert}_{i,L}(\mathit{division}(\texttt{order},\sigma))\\ \textbf{else}\ \texttt{assert}_{i,L}(\neg \textstyle\bigvee_{\sigma\in T_{Goals}} \mathit{division}(\texttt{order},\sigma))),\end{aligned}$$

where σ is a sequence of goals from a pre-given finite set of goals T_{Goals}. L waits a while before collecting the answers from G. Suppose two agents decide to respond: D and FA. D proposes the sequence σ_D:

$$\sigma_D = \langle \texttt{scan\_ruins}, \texttt{clear\_safe}, \texttt{clear\_risky}, \texttt{fist\_aid\_risky} \rangle.$$

In other words, first, scan the ruins of the building in search for survivors, then, clear the area with no sign of people, next, clear the area where survivors might be and finally, help survivors. In FA's view, the only difference is the goal order:

$$\sigma_{FA} = \langle \texttt{scan\_ruins}, \texttt{clear\_risky}, \texttt{fist\_aid\_risky}, \texttt{clear\_safe} \rangle.$$

As a response to *L's* call, the two agents utter:

$$\texttt{assert}_{D,L}(\mathit{division}(\texttt{order},\sigma_D)) \text{ and } \texttt{assert}_{FA,L}(\mathit{division}(\texttt{order},\sigma_{FA}))$$

These two assertions cause belief revision. The second step is voting. The preselected subset of answers collected previously (candidate terms) is $T_{\texttt{order},L} = \{\sigma_D, \sigma_{FA}\}$. L discloses this information to all other agents $i \in G$:

$$\texttt{assert}_{L,i}\left(\bigwedge_{t\in T_{\texttt{order},L}} \bigvee_{i\in G} \text{BEL}(i, \mathit{division}(\texttt{order},t))\right)$$

Subsequently, L opens voting on proposals by requests to all $i \in G$:

$$\begin{aligned}\texttt{request}_{L,i}(\ \bigwedge_{x,y\in T_{\texttt{order},L}}(\ \textbf{if}\ \ \text{BEL}(i, \mathit{division}(\texttt{order},x)) \wedge \text{prefer}(i,x,y)\\ \textbf{then}\ \ \texttt{assert}_{i,L}(\text{prefer}(i,x,y))))\end{aligned}$$

In step 3 (confirming), L announces that for example σ_{FA} won and calls potential opponents to start a persuasion dialogue, by sending a request to all other $i \in G$:

$$\texttt{request}_{L,i}(\ \textbf{if}\ \mathrm{BEL}(i, \neg \mathit{division}(\texttt{order}, \sigma_{FA})) \vee \bigvee_{t \in T_{\texttt{order},L}} \mathrm{prefer}(i, t, \sigma_{FA})\ \textbf{then}\ \texttt{assert}_{i,L}(\neg \mathit{division}(\texttt{order}, \sigma_{FA}) \vee \mathrm{prefer}(i, t, \sigma_{FA})))$$

If agent D prefers its own proposal, it raises an objection:

$$\texttt{assert}_{D,L}(\mathrm{prefer}(i, \sigma_D, \sigma_{FA}))$$

This is followed by L's challenge to provide a proof. At this point, the dialogue switches to persuasion, which has been discussed in [4, 5]. Step 4 (closure) follows the same pattern, leading, if successful, to a subgoal sequence σ such that $\mathit{division}(\texttt{order}, \sigma)$ holds.

The next stage is *means-end-analysis*, when every subgoal must be assigned a (complex) action realizing it. If the whole process concerning all subgoals from σ succeeds, there is an action sequence τ such that $\mathit{means}(\sigma, \tau)$ holds. The last phase, *action allocation*, results, if successful, in a plan P for which $\mathit{allocation}(\tau, P)$ holds. Finally, $\mathit{constitute}(\texttt{order}, P)$ is reached and planning terminates. There is now a basis to establish a collective commitment and to start working.

Although deliberation during teamwork is a complex process, all its phases can be naturally specified in TeamLog. First, central to teamwork theory, collective group attitudes are defined in terms of other informational and motivational attitudes (via fixpoint definitions). Then, the dynamic component allows specifying consequences of various speech acts, plans, and complex actions. These can be further applied as building blocks of different dialogues. The entire multimodal framework constituting TeamLog is presented in the recent book [5].

8 Conclusions and Future Work

We have introduced a novel approach to modeling deliberation dialogues in teamwork. Although dialogues and speech acts have been frequently used to model communication in multi-agent systems [8], the TeamLog solution is unique. The proposed scenario consists of four stages, during which agents submit their proposals, vote on preferred ones and challenge or concede the choice of the selected one. Depending on the context and aim of the system, a system designer decides on tactical aspects such as waiting time for the leader during answer collection, and manner of counting votes, possibly using weights depending on the agents and/or types of proposals.

The scenario specifies precisely when to embed other dialogues types into deliberation, as opposed to [11]. Different types of dialogues are strictly distinguished, and the boundary between them is clearly outlined, providing an appropriate amount of flexibility, enabling smooth teamwork (about the importance of dialogue embedding, see also [6]). In the course of deliberation, a social plan leading to the overall goal is created, belief revision is done and growth of knowledge can be observed. Finally, along with existing schemas for persuasion and information seeking [4], the most vital aspects of communication in TeamLog are now addressed.

In future, communication involving uncertain and possibly inconsistent information will be investigated, most probably requiring a new model of TeamLog.

9 Acknowledgements

This research has been supported by the Polish grant N N206 399334 and Dutch Vici grant NWO-277-80-001.

References

1. Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic*, 50(2):510–530, 1985.
2. Katie Atkinson, Trevor Bench-Capon, and Peter McBurney. Computational representation of practical argument. *Synthese*, 152:157–206, 2005.
3. J. L. Austin. *How to Do Things with Words*. Clarendon Press, Oxford, second edition, 1975. Edited by J. O. Urmson and M. Sbisa.
4. F. Dignum, B. Dunin-Kęplicz, and R. Verbrugge. Creating collective intention through dialogue. *Logic Journal of the IGPL*, 9:145–158, 2001.
5. B. Dunin-Kęplicz and R. Verbrugge. *Teamwork in Multi-Agent Systems: A Formal Approach*. Wiley, Chichester, 2010.
6. B. Dunin-Kęplicz and R. Verbrugge. Dialogue in teamwork. In *Proceedings of the 10th ISPE International Conference on Concurrent Engineering: Research and Applications*, pages 121–128, Rotterdam, 2003. A.A. Balkema.
7. M. Dziubiński, R. Verbrugge, and B. Dunin-Kęplicz. Complexity issues in multiagent logics. *Fundamenta Informaticae*, 75(1-4):239–262, 2007.
8. FIPA. 2002. `http://www.fipa.org/`.
9. Frank Guerin and Jeremy Pitt. Denotational semantics for agent communication language. In *AGENTS '01: Proceedings of the Fifth International Conference on Autonomous Agents*, pages 497–504, New York, NY, USA, 2001. ACM.
10. D. Harel, D. Kozen, and J. Tiuryn. *Dynamic Logic*. MIT Press, Cambridge, MA, 2000.
11. P. McBurney, D. Hitchcock, and S. Parsons. The eightfold way of deliberation dialogue. *International Journal of Intelligent Systems*, 22(1):95–132, 2007.
12. Peter McBurney and Simon Parsons. A denotational semantics for deliberation dialogues. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 86–93, Washington, DC, USA, 2004. IEEE Computer Society.
13. Simon Parsons and Peter McBurney. Argumentation-based dialogues for agent coordination. *Group Decision and Negotiation*, 12:415–439, 2003.
14. Henry Prakken. Formal systems for persuasion dialogue. *The Knowledge Engineering Review*, 21(2):163–188, 2006.
15. J.R. Searle and D. Vanderveken. *Foundations of Illocutionary Logic*. Cambridge University Press, Cambridge, 1985.
16. D.N. Walton and E.C.W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. State University of New York Press, Albany (NY), 1995.
17. M Wooldridge and NR Jennings. The cooperative problem-solving process. *Journal of Logic and Computation*, 9(4):563–592, 1999.

A Appendix: Axiom Systems

All axiom systems introduced here are based on the finite set $\mathcal{A}$ of n agents.

General Axiom and Rule The following cover propositional reasoning:
P1 All instances of propositional tautologies;
PR1 From φ and $\varphi \to \psi$, derive ψ;

Axioms and Rules for Individual Belief, Goal and Intention For beliefs, the KD45_n system for n agents is adopted, for intentions, KD_n, for goals, K_n.

Interdependencies Between Intentions and Other Attitudes For each $i \in \mathcal{A}$:
$\mathbf{A7}_{DB}$ $\text{GOAL}(i, \varphi) \to \text{BEL}(i, \text{GOAL}(i, \varphi))$
$\mathbf{A7}_{IB}$ $\text{INT}(i, \varphi) \to \text{BEL}(i, \text{INT}(i, \varphi))$
$\mathbf{A8}_{DB}$ $\neg\text{GOAL}(i, \varphi) \to \text{BEL}(i, \neg\text{GOAL}(i, \varphi))$
$\mathbf{A8}_{IB}$ $\neg\text{INT}(i, \varphi) \to \text{BEL}(i, \neg\text{INT}(i, \varphi))$
$\mathbf{A9}_{ID}$ $\text{INT}(i, \varphi) \to \text{GOAL}(i, \varphi)$

Axioms and Rule For General ("Everyone") and Common Belief
C1 $\text{E-BEL}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{BEL}(i, \varphi)$
C2 $\text{C-BEL}_G(\varphi) \leftrightarrow \text{E-BEL}_G(\varphi \wedge \text{C-BEL}_G(\varphi))$
RC1 From $\varphi \to \text{E-BEL}_G(\psi \wedge \varphi)$ infer $\varphi \to \text{C-BEL}_G(\psi)$

Axioms and Rule for Mutual and Collective Intentions
M1 $\text{E-INT}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{INT}(i, \varphi)$
M2 $\text{M-INT}_G(\varphi) \leftrightarrow \text{E-INT}_G(\varphi \wedge \text{M-INT}_G(\varphi))$
M3 $\text{C-INT}_G(\varphi) \leftrightarrow \text{M-INT}_G(\varphi) \wedge \text{C-BEL}_G(\text{M-INT}_G(\varphi))$
RM1 From $\varphi \to \text{E-INT}_G(\psi \wedge \varphi)$ infer $\varphi \to \text{M-INT}_G(\psi)$

Defining Axiom for Social Commitment

$$\begin{aligned}\text{COMM}(i, j, \alpha) \leftrightarrow \text{INT}(i, \alpha) \wedge \text{GOAL}(j, \mathit{done}(i, \alpha)) \wedge \\ \mathit{awareness}_{\{i,j\}}(\text{INT}(i, \alpha) \wedge \text{GOAL}(j, \mathit{done}(i, \alpha)))\end{aligned}$$

Defining Axiom for Strong Collective Commitment

$$\begin{aligned}&\text{S-COMM}_{G,P}(\varphi) \leftrightarrow \text{C-INT}_G(\varphi) \wedge \\ &\mathit{constitute}(\varphi, P) \wedge \text{C-BEL}_G(\mathit{constitute}(\varphi, P)) \wedge \\ &\bigwedge_{\alpha \in P} \bigvee_{i,j \in G} \text{COMM}(i, j, \alpha) \wedge \text{C-BEL}_G(\bigwedge_{\alpha \in P} \bigvee_{i,j \in G} \text{COMM}(i, j, \alpha))\end{aligned}$$

TeamLog denotes the union of the axioms for individual attitudes with the above axioms and rules for general and common beliefs and for general, mutual and collective intentions. $\text{TeamLog}^{\text{com}}$ denotes the union of TeamLog with the axioms for social and collective commitments (see [5, Chapter 4]). By $\text{TeamLog}^{\text{dyn}}$ we denote the union of $\text{TeamLog}^{\text{com}}$ with the axioms for dynamic operators, adopted from [10].