

Modeling Deliberation in Teamwork

Barbara Dunin-Keplicz

Institute of Informatics
Warsaw University,
Institute of Computer Science
Polish Academy of Sciences
Warsaw, Poland
keplicz@mimuw.edu.pl

Alina Strachocka

Institute of Informatics
Warsaw University
Warsaw, Poland
astrachocka@mimuw.edu.pl

Rineke Verbrugge

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. Taking off from the dialogue theory of Walton and Krabbe, the paper focuses on deliberation, the main type of dialogue during team planning. We provide a four-stage schema of deliberation dialogue along with semantics of adequate speech acts, filling the gap in logical modeling of communication during planning.

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 (Dunin-Keplicz and Verbrugge 2010), 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.

The paper is organized as follows. First, a brief introduction of dialogue and speech acts theory is given followed by discussion of a new four-stage model of deliberation dialogue. The next section provides explanations about teamwork and elaborates on the planning phase. Finally, conclusions and plans for future work are presented.

Speech Acts and Dialogues

Communication in MAS has two pillars: Walton and Krabbe's semi-formal theory of dialogue (1995) and the speech acts theory of Austin and Searle (1985; 1975). 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 in some situation. 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 (1975), that some utterances cannot be verified as true or false, led to the division of speech acts into *constatives*, that 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., suggesting, stating), *directives*, which get the hearer to do something (e.g., questioning), *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 (FIPA 2002). There have been many approaches to defining the semantics of speech acts (Parsons and McBurney 2003; Atkinson, Bench-Capon, and McBurney 2005; McBurney and Parsons 2004; Guerin and Pitt 2001; Singh 1993). Within the most popular *mentalist* approach, reflected in communication languages such as KQML and FIPA ACL (2002), speech acts are defined through their impact on agents' mental attitudes. On the other hand, some researchers view speech acts as primitive notions (Prakken 2006).

Semantics of Speech Acts

In TEAMLOG, deliberation dialogues are modeled via elementary speech acts *assert*, *concede* and *request*, and the compound speech acts *challenge* and *announce*, defined in terms of dynamic logic and described before in (Dignum, Dunin-Keplicz, and Verbrugge 2001). These speech acts are treated as ordinary actions and distinguished by their consequences. Utterances of speech acts often necessitate participants' belief revision, which may be handled by diverse methods, starting from Alchourrón, Gärdenfors, and Makinson.

In the sequel, the construction "if φ then α else β " will be used to abbreviate the dynamic logic expression $(\text{confirm}(\varphi); \alpha) \cup (\text{confirm}(\neg\varphi); \beta)$, and analogically for

“if φ then α ”, where $\text{confirm}(\varphi)$ refers to testing whether φ holds (see (Dunin-Kępicz and Verbrugge 2010, Chapter 6)). $[\beta]\varphi$ means that after performing β , φ holds.

$\text{BEL}(i, \varphi)$	agent i believes that φ
$\text{C-BEL}_G(\varphi)$	group G has the common belief that φ
$[\beta](\varphi)$	after performing social action β , φ holds
$\text{do-ac}(i, \alpha)$	agent i is just about to perform action α
$\text{division}(\varphi, \sigma)$	σ is the sequence of subgoals resulting from decomposition of φ
$\text{means}(\sigma, \tau)$	τ is the sequence of actions resulting from means-end analysis on sequence σ
$\text{allocation}(\tau, P)$	P is a social plan resulting from allocating the actions from τ to team members
$\text{constitute}(\varphi, P)$	P is a correct social plan for achieving φ
$\text{confirm}(\varphi)$	plan to test if φ holds at the given world
$\text{prefer}(i, x, y)$	agent i prefers x to y

Table 1: Formulas and their intended meaning

Consequences of assertions. $\text{assert}_{a,i}(\varphi)$ stands for agent a telling agent i that φ holds.

Definition 1 *The consequences of assertions:*

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

According to the fundamental assumption that agents are as truthful as they can be, each $\text{assert}(\varphi)$ obliges the sender to believe in φ . The recipient has two possibilities to react. Unless having beliefs conflicting with φ , it answers with a $\text{concede}_{i,a}$. Otherwise, with a $\text{challenge}_{i,a}$:

$$\begin{aligned} \neg \text{BEL}(i, \neg \varphi) &\rightarrow \text{do-ac}(i, \text{concede}_{i,a}(\varphi)), \\ \text{BEL}(i, \neg \varphi) &\rightarrow \text{do-ac}(i, \text{challenge}_{i,a}(\varphi)). \end{aligned}$$

Consequences of requests. $\text{request}_{a,i}(\alpha)$ stands for agent a requesting agent i to perform the action α . After a request for information about φ ($\alpha = \text{assert}_{i,a}(\varphi)$), the sender 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 :

$$\text{assert}_{i,a}(\neg(\text{BEL}(i, \varphi) \wedge \neg \text{BEL}(i, \neg \varphi))).$$

4. Either to assert that φ is the case or that it is not:
 $\text{BEL}(i, \varphi) \rightarrow \text{do-ac}(i, \text{assert}_{i,a}(\varphi))$, or
 $\text{BEL}(i, \neg \varphi) \rightarrow \text{do-ac}(i, \text{assert}_{i,a}(\neg \varphi))$.

The consequences are the same as for proper assertions.

Consequences of concessions. $\text{concede}_{a,i}(\varphi)$ stands for agent a 's communicating its positive attitude towards φ to i .

Definition 2 *The consequences of concessions:*

CCO $[\text{concede}_{a,i}(\varphi)] \text{BEL}(i, \text{BEL}(a, \varphi))$.

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.

Consequences of challenges. $\text{challenge}_{a,i}(\varphi)$ 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 3 *If $\varphi, \text{PROOF}(\varphi) \in \mathcal{L}$, $a, i \in \mathcal{A}$, then*

CH $\text{challenge}_{a,i}(\varphi) \equiv \text{assert}_{a,i}(\neg \varphi);$
 $\text{request}_{a,i}(\text{assert}_{i,a}(\text{PROOF}(\varphi)))$

The answer to the request in challenge should comply with the rules for information seeking. If i is able to prove φ , it should answer with $\text{assert}_{i,a}(\text{PROOF}(\varphi))$ being committed to $\text{PROOF}(\varphi)$. In return, a should refer to i 's previous answer. Thus¹, the consequences of challenge depend on the outcome of the dialogue and can be twofold.

Definition 4 *The consequences of challenges:*

CH1 $[\text{challenge}_{a,i}(\varphi)] (\text{BEL}(a, \varphi) \wedge \text{BEL}(i, \text{BEL}(a, \varphi))$
 $\wedge \text{BEL}(a, \text{PROOF}(\varphi)) \wedge \text{BEL}(i, \text{BEL}(a, \text{PROOF}(\varphi))))$

CH2 $[\text{challenge}_{a,i}(\varphi)] (\neg \text{BEL}(i, \varphi) \wedge \text{BEL}(a, \neg \text{BEL}(i, \varphi))$
 $\wedge \neg \text{BEL}(i, \text{PROOF}(\varphi)) \wedge \text{BEL}(a, \neg \text{BEL}(i, \text{PROOF}(\varphi))))$

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 $\text{announce}_{a,G}(\varphi)$ can be seen as a complex assertion standing for “agent a announces to group G that φ holds”. In addition to informing about φ , the agent passes a message that the same information has been delivered to the whole group. The group becomes commonly aware that φ .

Definition 5 *Consequences of announcements:*

CAN $[\text{announce}_{a,G}(\varphi)] \text{C-BEL}_G(\varphi)$.

Four-stage Model of Deliberation

The schema for deliberation dialogues presented below benefits from the model of McBurney, Hitchcock and Parsons (2007). 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 “ $\psi(x)$ ”² 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 (compare with the reconfiguration algorithm (Dunin-Kępicz and Verbrugge 2010)).

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

$$\begin{aligned} \text{request}_{a,i}(\text{if } \bigvee_{t \in T_\psi} \psi(t) \text{ then } \text{assert}_{i,a}(\psi(t)) \\ \text{else } \text{assert}_{i,a}(\neg \bigvee_{t \in T_\psi} \psi(t))). \end{aligned}$$

Agents have four ways of answering. 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:

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

Next, agent a opens the voting by a request to all $i \in G$:

$$\begin{aligned} \text{request}_{a,i}(\bigwedge_{x,y \in T_{\psi,a}} (\text{if } \psi(x) \wedge \text{prefer}(i, x, y) \text{ then} \\ \text{assert}_{i,a}(\text{prefer}(i, x, y)))) \end{aligned}$$

¹ Assuming the rule $\text{BEL}(a, \text{PROOF}(\varphi)) \rightarrow \text{BEL}(a, \varphi)$

² ψ is usually ungrounded, e.g. $\psi(x) = \text{president}(x)$. The answers are (partially) grounded terms, e.g., $\text{president}(\text{JohnSmith})$.

Again, the agents have four answering possibilities. 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, e.g., weighted by trust towards certain agents.

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

$$\text{request}_{a,i}(\text{if BEL}(i, \neg\psi(w)) \vee \bigvee_{t \in T_{\psi,a}} (\text{prefer}(i, t, w)))$$

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

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 $i \in G$:

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

In this situation, the remaining agents may concede:

$$\text{concede}_{i,a}(\neg\psi(w) \vee \text{prefer}(a, t, z))$$

or they may challenge the thesis:

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

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 to the final stage.

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

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. Let us place deliberation in the context of teamwork.

Stages of Teamwork

Teamwork, as the pinnacle of cooperation, is essential in multi-agent systems. The common division of teamwork into four stages originates from (Wooldridge and Jennings 1999), while a complete model, binding these stages to formalized team attitudes, can be found in (Dunin-Kępicz and Verbrugge 2010, Chapters 5 and 6). 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 certain common 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.
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 three phases: *task division*, leading to $\text{division}(\varphi, \sigma)$, (σ is a sequence of subgoals resulting from φ); *means-end analysis*, leading to $\text{means}(\sigma, \tau)$ (τ is the sequence of actions resulting from σ); and *action allocation*, leading to $\text{allocation}(\tau, P)$ (P is a social plan resulting from allocating τ to agents).

Success of the sequence of three phases is summed up by $\text{constitute}(\varphi, P)$ (P is a correct plan for achieving φ).

4. During **team action** agents execute their actions from the plan. However, in real situations, many actions are at risk of failure, calling for a necessary reconfiguration (Dunin-Kępicz and Verbrugge 2010), describing agents' behavior to salvage their goal in terms of intelligent replanning.

With each stage of teamwork, adequate notions and complex definitions in TEAMLOG are connected. There is no room for discussing them in detail, but see the book (Dunin-Kępicz and Verbrugge 2010).

Unveiling the Plan Formation Stage

The (formal) aim of plan formation is transition from collective intention to collective commitment, achieved by means of dialogue. Consider, as an example, a team of various unmanned winter service vehicles: snow plow (SP), snow blower (SB), salt spreader (SS), 10 lightweight robots with snow shovels ($LRS_1 \dots LRS_{10}$), transporter truck (TT), helicopter (H) and team leader (or initiator) with a goal to remove the snow from the roof: $\varphi = \text{office_SR}$. Suppose that 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 sequence of subgoals. *Leader* opens the deliberation dialogue by requesting all other $i \in G$ to share their ideas:

$$\text{request}_{\text{leader},i}(\text{if } \bigvee_{\sigma \in T_{Goals}} \text{division}(\text{office_SR}, \sigma))$$

$$\text{then assert}(\text{division}(\text{office_SR}, \sigma))$$

$$\text{else assert}(\neg \bigvee_{\alpha \in T_{Goals}} \text{division}(\text{office_SR}, \alpha)),$$

where σ is a sequence of goals from a pre-given finite set of goals T_{Goals} . The *leader* waits a while before collecting the answers from G . Suppose two agents decide to respond: SP and H . SP proposes the sequence σ_{SP} :

$$\sigma_{SP} = \langle \text{send_LRS_roof}, \text{clean_area}, \text{clean_roof}, \text{return} \rangle.$$

In other words, it proposes that first the LRS agents must be delivered to the roof, after that the building's surroundings must be cleared of snow, next the roof must be cleaned, then the team may return to the base. H proposes σ_H :

$$\sigma_H = \langle \text{send_LRS_roof}, \text{clean_roof}, \text{clean_area}, \text{return} \rangle.$$

In its view, the only difference is the ordering of goals. As a response to the *leader's* call, the two agents utter:

$$\text{assert}_{SP,\text{leader}}(\text{division}(\text{office_SR}, \sigma_{SP})) \text{ and}$$

$$\text{assert}_{H,\text{leader}}(\text{division}(\text{office_SR}, \sigma_H))$$

The consequence of these two assertions is belief revision. The second step is voting. *Leader* announces a pre-selected subset of answers collected in the previous step. The pre-selected set of candidate terms is $T_{\text{office_SR},\text{leader}} = \{\sigma_{SP}, \sigma_H\}$. Next comes disclosing this information to all other agents $i \in G$:

$$\text{assert}_{\text{leader},i} \left(\bigwedge_{t \in T_{\text{office_SR},\text{leader}}} \bigvee_{i \in G} \text{BEL}(i, \text{division}(\text{office_SR}, t)) \right)$$

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

$$\text{request}_{\text{leader},i} \left(\bigwedge_{x,y \in T_{\text{office_SR},\text{leader}}} (\text{if BEL}(i, \text{division}(\text{office_SR}, x)) \wedge \text{prefer}(i, x, y) \text{ then assert}_{i,\text{leader}}(\text{prefer}(i, x, y))) \right)$$

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

$\text{request}_{\text{leader},i}(\text{if BEL}(i, \neg \text{division}(\text{office_SR}, \sigma_{SP})) \vee$
 $\quad \vee_{t \in T_{\text{office_SR}}} \text{prefer}(i, t, \sigma_{SP}) \text{ then}$
 $\text{assert}_{i,\text{leader}}(\neg \text{division}(\text{office_SR}, \sigma_{SP}) \vee \text{prefer}(i, t, \sigma_{SP}))$)

If agent H prefers its own proposal, it raises an objection:
 $\text{assert}_{H,\text{leader}}(\text{prefer}(i, \sigma_H, \sigma_{SP}))$

This is followed by *leader's* challenge to provide a proof. At this point, the dialogue switches to persuasion, which has been discussed in (Dignum, Dunin-Kęplicz, and Verbrugge 2001; Dunin-Kęplicz and Verbrugge 2010). Step 4 (closure) follows the same pattern, leading, if successful, to a subgoal sequence σ such that $\text{division}(\text{office_SR}, \sigma)$ holds.

The next step 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 $\text{means}(\sigma, \tau)$ holds. The final step is action allocation, resulting, if successful, in a plan P for which $\text{allocation}(\tau, P)$ holds. Finally, $\text{constitute}(\text{office_SR}, P)$ is reached and planning terminates. There is now a basis to establish a collective commitment and to start working.

Although deliberation in the course of teamwork is a complex process, all its phases can be naturally specified in TEAMLOG. First, the central notions in a theory of teamwork, collective group attitudes, are defined in terms of other informational and motivational attitudes (via fixpoint definitions). Then, the dynamic component of TEAMLOG allows one to specify consequences of various speech acts, plans, and complex actions. These can be further applied as building blocks of different dialogues types. The entire multi-modal framework constituting TEAMLOG is presented in the recent book (Dunin-Kęplicz and Verbrugge 2010).

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 (FIPA 2002), 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. If other types of dialogues, like persuasion, are to be embedded into deliberation, this is precisely stated in the scenario, as opposed to (McBurney, Hitchcock, and Parsons 2007). The possibility to embed, for example, a negotiation within a persuasion in the confirming stage of deliberation about action allocation, provides an appropriate amount of flexibility, enabling smooth teamwork (about the importance of dialogue embedding, see also (Dunin-Kęplicz and Verbrugge 2003)). Different types of dialogues are strictly distinguished, and the boundary between them is clearly outlined. 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 (Dignum, Dunin-Kęplicz, and Verbrugge 2001), the most vital aspects of communication in TEAMLOG are now addressed.

In future, communication in the presence of uncertain and possibly inconsistent information will be investigated, most probably requiring a new model of TEAMLOG.

Acknowledgements

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

References

- Alchourrón, C. E.; Gärdenfors, P.; and Makinson, D. 1985. On the logic of theory change: Partial meet contraction and revision functions. *Journal of Symbolic Logic* 50(2):510–530.
- Atkinson, K.; Bench-Capon, T.; and McBurney, P. 2005. Computational representation of practical argument. *Synthese* 152:157–206.
- Austin, J. L. 1975. *How to Do Things with Words*. Oxford: Clarendon Press, second edition. Edited by J. O. Urmson and M. Sbisà.
- Dignum, F.; Dunin-Kęplicz, B.; and Verbrugge, R. 2001. Creating collective intention through dialogue. *Logic Journal of the IGPL* 9:145–158.
- Dunin-Kęplicz, B., and Verbrugge, R. 2010. *Teamwork in Multi-Agent Systems: A Formal Approach*. Chichester: Wiley.
- Dunin-Kęplicz, B., and Verbrugge, R. 2003. Dialogue in teamwork. In *Proceedings of the 10th ISPE International Conference on Concurrent Engineering: Research and Applications*, 121–128. Rotterdam: A.A. Balkema.
- FIPA. 2002. <http://www.fipa.org/>.
- Guerin, F., and Pitt, J. 2001. Denotational semantics for agent communication language. In *AGENTS '01: Proceedings of the Fifth International Conference on Autonomous Agents*, 497–504. New York, NY, USA: ACM.
- McBurney, P., and Parsons, S. 2004. A denotational semantics for deliberation dialogues. In *AAMAS '04: Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, 86–93. Washington, DC, USA: IEEE Computer Society.
- McBurney, P.; Hitchcock, D.; and Parsons, S. 2007. The eightfold way of deliberation dialogue. *International Journal of Intelligent Systems* 22(1):95–132.
- Parsons, S., and McBurney, P. 2003. Argumentation-based dialogues for agent coordination. *Group Decision and Negotiation* 12:415–439.
- Prakken, H. 2006. Formal systems for persuasion dialogue. *The Knowledge Engineering Review* 21(2):163–188.
- Searle, J., and Vanderveken, D. 1985. *Foundations of Illocutionary Logic*. Cambridge: Cambridge University Press.
- Singh, M. 1993. A semantics for speech acts. *Annals of Mathematics and Artificial Intelligence* 8:47–71.
- Walton, D., and Krabbe, E. 1995. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning*. Albany (NY): State University of New York Press.
- Wooldridge, M., and Jennings, N. 1999. The cooperative problem-solving process. *Journal of Logic and Computation* 9(4):563–592.