

# Collective attitudes for teamwork in dynamic environments

Barbara Dunin-Kęplicz<sup>1,2</sup> and Rineke Verbrugge<sup>3</sup>

<sup>1</sup> Institute of Informatics, Warsaw University, Banacha 2, 02-097 Warsaw, Poland

<sup>2</sup> Institute of Computer Science, Polish Academy of Sciences, Ordona 21, 01-237 Warsaw, Poland, [keplicz@mimuw.edu.pl](mailto:keplicz@mimuw.edu.pl)

<sup>3</sup> Department of Artificial Intelligence, University of Groningen, Grote Kruisstraat 2/1, 9712 TS Groningen, The Netherlands, [rineke@ai.rug.nl](mailto:rineke@ai.rug.nl)

**Abstract.** This article presents a theory of teamwork that may be applicable for example in multiagent systems. This formal model aims to enhance human-computer cooperation by precisely defining the notions necessary to methodologically organize the work of a cooperative team. The framework consists of both *a static part*, defining collective motivational attitudes in such a way that the system developer can adapt them to the circumstances, and *a dynamic part* monitoring the changes in team attitudes during the course of cooperative problem solving (CPS).

In this article we focus on the static part of the theory, including its cognitive aspects. First, the notion of *collective intention* in teams of agents is investigated. Starting from individual *intentions*, *goals*, and *beliefs* defining agents' local asocial attitudes, we arrive at an understanding of collective intention in cooperative teams as a rather strong concept: it implies that all members *intend* for all others to share that intention. This way a team is glued together by collective intention, and exists as long as this attitude holds, after which it may disintegrate.

Collective intentions are formalized in a multi-modal logical framework. Together with individual and collective knowledge and/or belief, collective intention constitutes a basis for preparing a plan, reflected in the strongest attitude, i.e., in *collective commitment*, defined and investigated in the next part. Different versions of collective commitments that are applicable in different situations, differ with respect to the *aspects* of teamwork of which the agents involved are aware, and the *kind* of awareness present within a team. This way a kind of tuning mechanism is provided for the system developer to tune a version of collective commitment fitting the circumstances. Finally, a few exemplar versions of collective commitment resulting from instantiating the general tuning scheme are presented.

This paper is meant as a survey of the teamwork theory that has been presented in a much more technical manner in [1, 2].

## 1 Introduction

In *multiagent systems* (MAS) one of the central issues is the study of how groups work, and how the technology enhancing group interaction

can be implemented. From the distributed Artificial Intelligence perspective, multiagent systems are computational systems in which a collection of loosely-coupled autonomous agents interact in order to solve a given problem. As this problem is usually beyond the agents' individual capabilities, agents exploit their ability to *communicate*, *cooperate*, *coordinate*, and *negotiate* with one another. Apparently, the type of social interactions involved depends on circumstances and may vary from altruistic cooperation through to open conflict. A paradigmatic example of joint activity is *cooperative problem solving* (CPS) in which a group of autonomous agents choose to work together, both in advancement of their own goals as well as for the good of the system as a whole.

### 1.1 Intentional multiagent systems

Some MAS are referred to as *intentional systems*. In such systems, in order to give a representation of the mental states and cognitive processes involved in a multiagent system, agents are represented as maintaining an intentional stance towards their environment. Such systems realize the *practical reasoning* paradigm ([3]) – the process of deciding, moment by moment, which action to perform in the furtherance of our goals. The best known and most influential are *belief-desire-intention systems*. BDI-agents are characterized by a “mental state” described in terms of *beliefs*, corresponding to the information the agent has about the environment; *desires*, representing options available to the agent, i.e. different states of affairs that the agent may choose to commit to; and *intentions* representing the chosen options. Ultimately, in our approach, intentions are viewed as an inspiration for a goal-directed activity, reflected in commitments. While beliefs are viewed as the agent's *informational* attitudes, desires or goals, intentions, and commitments refer to its *motivational* attitudes. These collective notions are considered in the context of teams of agents defined in [4]:

A *team* is a group in which the agents are restricted to having a common goal of some sort. Typically, team-members cooperate and assist each other in achieving their common goal.

*Collective intention*, as a specific joint mental attitude, is the central topic addressed in *teamwork*. We agree with [5] that:

Joint intention by a team does not consist merely of simultaneous and coordinated individual actions; to act together, a team must

be aware of and care about the status of the group effort as a whole.

In fact, we assume that a team is constituted as soon as a collective intention is present among the members, and stays together as long as the collective intention persists.

We agree with Bratman that in human practical reasoning, intentions are first class citizens, in the sense that they are not reducible to beliefs and desires [3]. They form a rather special consistent subset of an agent's goals, that the agent wants to focus on for the time being. Thus they create a screen of admissibility for the agent's further, possibly long-term, deliberation. In contrast to [6] we are interested in generic characteristics of intentions, resigning from classifying them further along different dimensions. The most crucial aspect of intentions is that they are considered as an inspiration for goal-directed activity, reflected in the strongest motivational attitudes, that is in social (bilateral) and collective commitments. In our view, social commitments are related to individual actions, while collective commitments are related to plan-based team actions. The essential characteristics of both types of commitments is that they directly lead to action execution. Commitments will be treated in depth in section 4.

In our work on collective intentions, the objective is two-fold. First, we aim to formally characterize what it means for a team to have a collective intention towards a common goal (represented for example as a state of the world to be achieved). This characterization will be done using multi-modal logic, in the form of a *static* theory, comprising a descriptive view on collective intentions.

Secondly, let us remember that a team of agents in a multiagent system operates in a dynamic and often unpredictable environment. Such an environment poses the problem that team members may fail to bring their tasks to a good end or new opportunities may appear. We treated this *reconfiguration problem* in [7, 8], where collective intentions are maintained by properly adapting collective commitments to the changing circumstances. This contributes to the *dynamic*, more prescriptive theory of collective intentions. We shall not treat this dynamic theory in detail in the present article.

Section 3 falls squarely in the scope of the objective to define a static theory of collective intentions. A definition of collective intention is presented, together with examples of situations in which it applies. The notion is formalized in a multi-modal logical framework. Completeness of this logic with respect to an appropriate class of Kripke models has been

proved in [1]. Thus, a Computational Logic framework for specifying MAS involved in CPS has been provided. It is assumed that presented definitions reflect solely vital aspects of motivational attitudes, leaving room for case-specific extensions. This makes the framework flexible and not overloaded.

## 1.2 The role of intentions in practical reasoning

Practical reasoning is the form of reasoning that is aimed at conduct rather than knowledge. The cycle of this reasoning involves:

1. repeatedly updating beliefs about the environment
2. deciding what options are available
3. “filtering” these options to determine new intentions
4. creating commitments on the basis of intentions
5. performing actions in accordance with commitments.

Practical reasoning involves two important processes: deciding *what* goals need to be achieved, and then *how* to achieve them. The former process is known as *deliberation*, the latter as *means-end-reasoning*. A key concept in the theory of practical reasoning is that of *intention*. In [9] it is characterized from the psychological viewpoint:

“The concept of intention was (and still is) one of the most controversial in the history of psychology. Certain people – the eliminativists – purely and simply refuse to introduce this concept into their theories, claiming not only that it is useless, but also that it mindlessly confuses the issues. Others, in contrast, think of it as one of the essential concepts of psychology and that it should be given a central role, for it constitutes a keystone of the explanation of human behaviour in terms of mental states. Finally, the psycho-analytical school sees it as merely a vague concept which is handy in certain cases, but which should generally be replaced by desire and drives, which alone are capable of taking account of the overall behaviour of the human being in his or her aspirations and suffering.”

We do not aim to present a psychologically sound theory of motivations driving human behaviour. We are interested in studying those motivational aspects that are involved in the *rational* decision making process. Thus, we disregard irrational drives and desires, that make human behaviour difficult to interpret and to predict. We do not consider

any specific notion of *rationality*, in particular the economic one used in game theory. The only assumption made here is that agents are logical reasoners.

There is common agreement that intentions play a number of important roles in practical reasoning [3, 6]:

- I1 *Intentions drive means-end-reasoning.*
- I2 *Intentions constrain future deliberation.*
- I3 *Intentions persist.*
- I4 *Intentions influence beliefs upon which future practical reasoning is based.*

A key problem in the design of BDI-agents is how to achieve a good balance between these different concerns. It becomes especially important when an agent needs to drop some of its intentions. This happens for many different reasons: because the intentions will never be achieved, they are achieved already or there are no longer reasons supporting them. Thus, from time to time an agent's intentions should be reconsidered. This leads to the problem of balancing *pro-active*, (i.e. goal-directed) and *reactive* (i.e. event-driven) behaviour. We try to maintain this balance very carefully on both the individual and the collective level. The problem of persistence of intentions is expressed in an agent's *intention strategies*, addressing the question: *when and how can an agent responsibly drop its intentions?* One answer to this question is discussed in [10]. The collective level is apparently more complex. Collective intention helps the team to monitor its behaviour during teamwork: even if some members drop their individual intentions, the team replans, aiming that the collective intention is ultimately realized. The precise description of this reconfiguration process can be found in [11, 7], where an agent's pro-activeness and reactivity are implicitly or explicitly involved in consecutive stages of the reconfiguration algorithm.

### 1.3 Calibrating team attitudes to the circumstances

Variety is the core of multiagent systems. This statement expresses the many dimensions on which MAS is distinguished from distributed AI. The basic assumption underlying MAS is relaxing the constraints that were fixed before, in order to meet the needs of goal-directed behaviour in a dynamic and unpredictable environment. This is reflected in complex and possibly flexible patterns of interaction in MAS. Together with autonomy of agents and social structure of cooperative groups this determines the novelty of the agent-based approach.

In this article we focus on collective motivational attitudes, i.e. intentions and commitments in CPS. What characterizes collective notions is an interplay between environmental and social aspects, which may become rather complex nowadays due to the increasing complexity of MAS. For example, when asking what it means for a group of agents to be *collectively committed* to do something, both the circumstances in which the group is acting and properties of the organization it is part of, have to be taken into account. This implies the importance of differentiating the scope and strength of the notion of collective commitment. The resulting characteristics may differ significantly, and even become logically incomparable.

The aim of section 4 is to formally model different aspects of collective commitments, including different scopes and degrees of awareness of cooperating agents. (The case of competition is not included.) The idea of a dial to be used to tune the nature of the commitment to the particular purpose seems to be both technically interesting and intuitively appealing. We intend to provide a sort of *tuning mechanism* which enables the system developer to *calibrate* a type of collective commitment fitting the circumstances, analogously to adjusting dials on a sound system. The appropriate dials, characterized in the sequel, belong to the device representing a general schema of collective commitment. The resulting notion of (group) commitment, described in multimodal logics, may then be naturally implemented in a created multiagent system. This way the tuning mechanism may be viewed as a bridge between theory and practice.

In order to illustrate the expressive power of such a sort of *tuning machine*, five definitions of commitments corresponding to different teamwork types occurring in practice are presented in subsection 4.3. Apparently, the entire spectrum of possibilities is much wider, due to the number of possibly independent choices to be made.

## 2 The logical background

As mentioned before, we propose the use of multi-modal logics to formalize agents' informational and motivational attitudes as well as actions they perform. A good introduction to modal and multi-modal logics can be found in [12].

As in this article we restrict ourselves to the static aspects of the agents' mental states, we present solely axioms relating agents' attitudes with respect to *propositions*, reflecting a particular state of affairs. Table

1 gives some important formulas appearing in this article, together with their intended meanings. The symbol  $\varphi$  denotes a proposition.

BEL( $a, \varphi$ )	agent $a$ has the belief that $\varphi$
E-BEL $_G(\varphi)$	every agent in group $G$ has the belief that $\varphi$
C-BEL $_G(\varphi)$	group $G$ has the common belief that $\varphi$
GOAL( $a, \varphi$ )	agent $a$ has as a goal to achieve $\varphi$
INT( $a, \varphi$ )	agent $a$ has the intention to achieve $\varphi$
E-INT $_G(\varphi)$	every agent in group $G$ has the individual intention to achieve $\varphi$
M-INT $_G(\varphi)$	group $G$ has the mutual intention to achieve $\varphi$
C-INT $_G(\varphi)$	group $G$ has the collective intention to achieve $\varphi$
C-COMM $_{G,P}(\varphi)$	group $G$ has a collective commitment to achieve $\varphi$ by plan $P$
R-COMM $_{G,P}(\varphi)$	group $G$ has robust collective commitment to achieve $\varphi$ by plan $P$
S-COMM $_{G,P}(\varphi)$	group $G$ has strong collective commitment to achieve $\varphi$ by plan $P$
W-COMM $_{G,P}(\varphi)$	group $G$ has weak collective commitment to achieve $\varphi$ by plan $P$
T-COMM $_{G,P}(\varphi)$	group $G$ has team commitment to achieve $\varphi$ by plan $P$
D-COMM $_{G,P}(\varphi)$	group $G$ has distributed commitment to achieve $\varphi$ by plan $P$

**Table 1.** Formulas and their intended meaning

## 2.1 The logical language

The language  $\mathcal{L}$  is based on a denumerable set  $\mathcal{P}$  of *propositional symbols*, and a finite set  $\mathcal{A}$  of *agents*, denoted by numerals  $1, 2, \dots, n$ , and a set  $\mathcal{At}$  of atomic actions. It is geared towards expressing the dynamics of attitudes in CPS and does not include temporal operators, for example.

### Definition 1 (Formulas).

We inductively define a set of formulas  $\mathcal{L}$  as follows.

- F1** each atomic proposition  $p \in \mathcal{P}$  is a formula;
- F2** if  $\varphi$  and  $\psi$  are formulas, then so are  $\neg\varphi$  and  $\varphi \wedge \psi$ ;
- F3** if  $\varphi$  is a formula,  $\alpha$  is an individual action,  $i, j \in \mathcal{A}$ ,  $G \subseteq \mathcal{A}$ , and  $P$  a social plan expression, then the following are formulas:
  - epistemic** BEL( $i, \varphi$ ), E-BEL $_G(\varphi)$ , C-BEL $_G(\varphi)$ ;
  - motivational** GOAL( $i, \varphi$ ), INT( $i, \varphi$ ), C-INT $_G(\varphi)$ , COMM( $i, j, \alpha$ ), C-COMM $_{G,P}(\varphi)$ , R-COMM $_{G,P}(\varphi)$ , S-COMM $_{G,P}(\varphi)$ , W-COMM $_{G,P}(\varphi)$ , T-COMM $_{G,P}(\varphi)$ ;

Social plans combine individual actions into a group action, and are represented by social plan expressions. The class of social plan expressions  $\mathcal{Sp}$  is based on a set of individual actions  $\mathcal{Ac}$  based on  $\mathcal{At}$ , which are application-dependent and will not be defined here. An example of a social plan is given below the definition.

**Definition 2 (Social plan expressions).**

The class  $\mathcal{Sp}$  of social plan expressions is defined inductively as follows:

- SP1** if  $\alpha \in \mathcal{Ac}$  and  $i \in \mathcal{A}$ , then  $\langle \alpha, i \rangle$  is a social plan expression;  
**SP2** If  $\alpha$  and  $\beta$  are social plan expressions, then  $(\alpha; \beta)$  (sequential composition) and  $(\alpha \parallel \beta)$  (parallelism) are social plan expressions.

Let us give a simple **example** of a social plan expression. Consider a team consisting of three agents  $s$  (syntax expert),  $m$  (semantics expert) and  $c$  (exam checker) who have as collective intention to compose a linguistics exam. Suppose during planning they define the following complex individual actions:  $coSyn$ ,  $coSem$  (to compose questions on syntax / semantics),  $chSyn$ ,  $chSem$  (to check questions on syntax / semantics). One possible social plan they can come up with is the following. First, the syntax expert and the semantics expert work in parallel composing questions about their fields, and subsequently the exam checker questions on correctness, first the syntax ones and then the semantics ones:

$$P = \langle \langle \langle \langle coSyn, s \rangle \parallel \langle coSem, m \rangle \rangle; \langle chSyn, c \rangle \rangle; \langle chSem, c \rangle \rangle.$$

**2.2 Axioms for beliefs**

To represent beliefs, we adopt the standard  $KD45_n$ -system for  $n$  agents as explained in [13], containing the following axioms and rules for agents  $i = 1, \dots, n$ :

- A1** All instantiations of tautologies of the propositional calculus  
**A2**  $BEL(i, \varphi) \wedge BEL(i, \varphi \rightarrow \psi) \rightarrow BEL(i, \psi)$  (Belief Distribution)  
**A4**  $BEL(i, \varphi) \rightarrow BEL(i, BEL(i, \varphi))$  (Positive Introspection)  
**A5**  $\neg BEL(i, \varphi) \rightarrow BEL(i, \neg BEL(i, \varphi))$  (Negative Introspection)  
**A6**  $\neg BEL(i, \perp)$  (Belief Consistency)  
**R1** From  $\varphi$  and  $\varphi \rightarrow \psi$  infer  $\psi$  (Modus Ponens)  
**R2** From  $\varphi$  infer  $BEL(i, \varphi)$  (Belief Generalization)

One can define modal operators for group belief. Let  $G \subseteq \{1, \dots, n\}$  be a group. The formula  $E-BEL_G(\varphi)$  is meant to stand for “every agent in group  $G$  believes  $\varphi$ ”, which corresponds to the following axiom:

$$\mathbf{C1} \quad E-BEL_G(\varphi) \leftrightarrow \bigwedge_{i \in G} BEL(i, \varphi)$$

A traditional way of lifting single-agent concepts to multiagent ones is through the use of *common belief*  $C-BEL_G(\varphi)$ . This rather strong operator is similar to the more usual one of common knowledge, except that a common belief among a group that  $\varphi$  need not imply that  $\varphi$  is true.



C-BEL<sub>G</sub>( $\varphi$ ) is meant to be true if everyone in  $G$  believes  $\varphi$ , and everyone in  $G$  believes that everyone in  $G$  believes  $\varphi$ , etc.

Let E-BEL<sub>G</sub><sup>2</sup>( $\varphi$ ) be an abbreviation for E-BEL<sub>G</sub>(E-BEL<sub>G</sub>( $\varphi$ )), and similarly for every  $k \geq 1$ , E-BEL<sub>G</sub><sup>k</sup>( $\varphi$ ) stands for abbreviation of E-BEL<sub>G</sub> applied  $k$  times to  $\varphi$ . Thus C-BEL<sub>G</sub>( $\varphi$ ) stands for the combination of E-BEL<sub>G</sub><sup>k</sup>( $\varphi$ ) for all  $k \geq 1$ . Unfortunately, the language only contains *finite* conjunctions. But we can still define common belief using some technical tricks (see [13] for an explanation), namely by the combination of an axiom and a rule:

**C2** C-BEL<sub>G</sub>( $\varphi$ )  $\leftrightarrow$  E-BEL<sub>G</sub>( $\varphi \wedge$  C-BEL<sub>G</sub>( $\varphi$ ))

**RC1** From  $\varphi \rightarrow$  E-BEL<sub>G</sub>( $\psi \wedge \varphi$ ) infer  $\varphi \rightarrow$  C-BEL<sub>G</sub>( $\psi$ ) (Induction Rule)

The resulting system is called  $KD45_n^C$ . In the sequel, we will use the following standard properties of C-BEL<sub>G</sub> (see for example [13, exercise 3.11]).

**Lemma 1.** *Let  $G \subseteq \{1, \dots, n\}$  be given. Then the following hold for all formulas  $\varphi, \psi$ :*

- C-BEL<sub>G</sub>( $\varphi \wedge \psi$ )  $\leftrightarrow$  C-BEL<sub>G</sub>( $\varphi$ )  $\wedge$  C-BEL<sub>G</sub>( $\psi$ )
- C-BEL<sub>G</sub>( $\varphi$ )  $\rightarrow$  C-BEL<sub>G</sub>(C-BEL<sub>G</sub>( $\varphi$ ))

A problem with standard modal logics for beliefs and knowledge is that agents are formalized as being logically omniscient: they believe all theorems, as well as all logical consequences of their beliefs. Logical omniscience definitely does not apply to human beings, who have only limited time available and have bounded rationality: it is unrealistic to assume that they believe every logical theorem, however complicated. There are several possible solutions to the problem, for example involving syntactic operators for awareness and explicit belief. Good references to the logical omniscience problem and its possible solutions are [14, Chapter 2] and [13, Chapter 9].

We acknowledge that these problems are important and should be adequately solved. Here we focus on the formalization of collective motivational attitudes needed for teamwork, and for the time being we choose to base it on the relatively simple, even if not practically adequate, logic for common belief defined above. Thus, we view common belief as a good *abstraction tool* to study teamwork.

**Degrees of belief in a group** It is well-known that for teamwork, as well as coordination, it often does not suffice that a group of agents all

believe a certain proposition (E-BEL<sub>G</sub>( $\psi$ )), but they should collectively believe it (C-BEL<sub>G</sub>( $\psi$ )).

One advantage of common belief is that if C-BEL<sub>G</sub> holds for  $\psi$ , then C-BEL<sub>G</sub> also holds for all logical consequences of  $\psi$ . Thus, agents reason in a similar way from  $\psi$  and collectively believe in this similar reasoning and the final conclusions. In short, one could say that common belief is hard to achieve, but easy to understand.

In cases in which only E-BEL<sub>G</sub>( $\psi$ ) has been established, it is much more difficult for agents to maintain a model of the other team members with respect to  $\psi$  and its consequences. However, establishing E-BEL<sub>G</sub>( $\psi$ ) places much less constraints on the communication medium than establishing C-BEL<sub>G</sub>( $\psi$ ). Thus, the system developer's decision about the level  $k$  of group belief (E-BEL<sub>G</sub> <sup>$k$</sup> ( $\psi$ )) to be established, hinges on determining a good balance between communication and reasoning taking into account a particular application.

### 2.3 Axioms for individual and social motivational attitudes

Our framework to describe motivational attitudes and related aspects is minimal in the sense that we aim to deal with concise necessary and sufficient conditions. Additional aspects appearing on the stage in specific cases may be addressed by refining the system and adding new axioms. This subsection focuses on individual goals and intentions, and gives a short overview of our choice of axioms (adapted from [10]). In this article, we leave out of consideration the aspects of time and action in order to focus on the main problem, the definition of collective intentions and commitments in terms of more basic attitudes.

For the motivational operators GOAL and INT the axioms include the basic modal system  $K$ , which we adapt for  $n$  agents to  $K_n$ . For  $i = 1, \dots, n$  the following axioms and rules are included:

- A1** All instantiations of tautologies of the propositional calculus
- R1** From  $\varphi$  and  $\varphi \rightarrow \psi$  infer  $\psi$  (Modus Ponens)
- A2<sub>G</sub>** GOAL( $i, \varphi$ )  $\wedge$  GOAL( $i, \varphi \rightarrow \psi$ )  $\rightarrow$  GOAL( $i, \psi$ ) (Goal Distribution)
- A2<sub>I</sub>** INT( $i, \varphi$ )  $\wedge$  INT( $i, \varphi \rightarrow \psi$ )  $\rightarrow$  INT( $i, \psi$ ) (Intention Distribution)
- R2<sub>G</sub>** From  $\varphi$  infer GOAL( $i, \varphi$ ) (Goal Generalization)
- R2<sub>I</sub>** From  $\varphi$  infer INT( $i, \varphi$ ) (Intention Generalization)

In a BDI system, an agent's activity starts from goals. In general, it may have many different objectives which will not all be pursued. As opposed to intentions, goals are not directly related to actions, so an agent

can behave rationally, even though it has different inconsistent goals. Thus, in contrast to Rao and Georgeff we adopt the basic system  $K_n$  for goals. Then, the agent chooses a limited number of these goals to be intentions. Here we do not discuss how intentions are formed from a set of goals (but see [15, 16]). In any case, we assume that intentions are chosen in such a way that consistency is preserved. Thus for intentions we assume, as Rao and Georgeff do, that they should be consistent. This can be formulated as follows:

**A6<sub>I</sub>**  $\neg\text{INT}(i, \perp)$  for  $i = 1, \dots, n$  (Intention Consistency Axiom)

## 2.4 Interdependencies between attitudes

Interdependencies between belief and individual motivational attitudes are expressed by the following axioms for  $i = 1, \dots, n$ :

**A7<sub>GB</sub>**  $\text{GOAL}(i, \varphi) \rightarrow \text{BEL}(i, \text{GOAL}(i, \varphi))$  (Positive Goal Introspection).

**A7<sub>IB</sub>**  $\text{INT}(i, \varphi) \rightarrow \text{BEL}(i, \text{INT}(i, \varphi))$  (Positive Intention Introspection).

**A8<sub>GB</sub>**  $\neg\text{GOAL}(i, \varphi) \rightarrow \text{BEL}(i, \neg\text{GOAL}(i, \varphi))$  (Negative Goal Introspection).

**A8<sub>IB</sub>**  $\neg\text{INT}(i, \varphi) \rightarrow \text{BEL}(i, \neg\text{INT}(i, \varphi))$  (Negative Intention Introspection).

These four axioms express that agents are aware of the goals and intentions they have, as well as of the lack of those that they do not have.

We assume that every intention corresponds to a goal:

**A9<sub>IG</sub>**  $\text{INT}(i, \varphi) \rightarrow \text{GOAL}(i, \varphi)$  (Intention implies goal)

This means that if an agent adopts a formula as an intention, it should have adopted that formula as a goal to achieve, which satisfies Bratman's notion that an agent's intentions form a specific subset of its goals [3]. Rao and Georgeff adopt this axiom as *goal-intention compatibility* for a restricted class of formulas [10].

Note that in our system it fortunately does not follow that an agent intends all the consequences it believes its intentions to have, i.e. the side-effects of its intentions. For a discussion of this "side-effect problem" for intentions, see [3, 6, 10].

Human beings reason under bounded rationality, due to cognitive limitations and time constraints. An overview from the point of view of cognitive psychology about ways in which motivations may influence beliefs in human reasoning, is given in the chapter Motivational Thinking in [17].

## 2.5 Social commitments

As Castelfranchi showed, it is important to distinguish between individual intentions, bilateral commitments, and collective motivational attitudes [18]. A social (bilateral) commitment is not as strong as a collective one, but stronger than an individual intention. If an agent commits to a second agent to do something ( $\text{COMM}(i, j, \varphi)$ ), then the first agent should have the *intention* to do that ( $\text{INT}(i, \varphi)$ ). Moreover, the second one should be *interested* in the first one fulfilling its intention (formalized as  $\text{GOAL}(j, \text{stit}(i, \varphi))$ : agent  $j$  has the goal that agent  $i$  sees to it that  $\varphi$  is achieved.). These two conditions (inspired by [18]), need to be enhanced by the condition expressing the agents' awareness about the situation, i.e. about their individual attitudes. Such awareness is usually expressed in terms of common belief. Here the defining axiom for social commitments is formulated with respect to propositions  $\varphi$  as well as to actions  $\alpha$ :

$$\text{COMM}(i, j, \varphi) \leftrightarrow \text{INT}(i, \varphi) \wedge \text{GOAL}(j, \text{stit}(i, \varphi)) \wedge$$

$$\text{C-BEL}_{\{i, j\}}(\text{INT}(i, \varphi) \wedge \text{GOAL}(j, \text{stit}(i, \varphi)))$$

where  $\text{stit}(i, \varphi)$  means that agent  $i$  sees to it (brings it about) that  $\varphi$  becomes true (see [19]).

$$\text{COMM}(i, j, \alpha) \leftrightarrow \text{INT}(i, \alpha) \wedge \text{GOAL}(j, \text{done}(i, \alpha)) \wedge$$

$$\text{C-BEL}_{\{i, j\}}(\text{INT}(i, \alpha) \wedge \text{GOAL}(j, \text{done}(i, \alpha)))$$

where  $\text{done}(i, \alpha)$  means that agent  $i$  has just executed action  $\alpha$ ; for a formal treatment of *done*, see [8].

## 3 Defining collective intention

In our approach, teams are created on the basis of *collective intentions*, which are defined in subsection 3.1 for the standard context and in [1] for less ideal circumstances. A team exists as long as the collective intention between team members exists. However here we abstract from the ways in which teams are formed, and refer the interested reader to [20, 21, 15, 22, 23].

In the sequel, collective intention is viewed as an inspiration for the creation of collective commitment leading directly to action execution. This is based on the linguistic tradition that intentions typically ultimately lead to actions, however, the immediate triggers of these actions are commitments.

In the philosophical and MAS literature there is an ongoing discussion as to whether collective intentions may be reduced to individual ones plus

common beliefs about them (see [18, 24, 25]). Even though our definition seems to be reductive, it involves nested intentions and collective epistemic operators, and for this reason is deeper than a simple compound built out of individual intentions and collective beliefs about them by propositional connectives only.

### 3.1 Collective intentions: the standard case

In this article we focus on strictly cooperative teams (see for example [26, 27] for good philosophical discussions of cooperation). This makes the definition of collective intention rather strong. In such teams, a necessary condition for a collective intention is that all members of the team  $G$  have the associated individual intention  $\text{INT}(i, \varphi)$  towards the state of the world represented by proposition  $\varphi$ . In fact, this condition is taken to give the full definition of collective intention in [28] (see [29] for a similar definition of collective goal). However, this is certainly not sufficient. Imagine that two agents want to achieve the same state of the world but are in a competition about this, willing to achieve it exclusively. Therefore, to exclude the case of competition, all agents should *intend* all members to have the associated individual intention, as well as the intention that all members have the individual intention, and so on; we call such a mutual intention  $\text{M-INT}_G(\varphi)$ . Furthermore, all members of the team are aware of this mutual intention, that is, they have a collective belief about this:  $\text{C-BEL}_G(\text{M-INT}_G(\varphi))$ . Of course, team members remain autonomous maintaining their other motivational attitudes, and may even be in competition about other issues.

In order to formalize the above two conditions,  $\text{E-INT}_G(\varphi)$  (standing for “everyone intends”) is defined by the following axiom:

$$\mathbf{M1} \quad \text{E-INT}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{INT}(i, \varphi).$$

The mutual intention  $\text{M-INT}_G(\varphi)$  is meant to be true if everyone in  $G$  intends  $\varphi$ , everyone in  $G$  intends that everyone in  $G$  intends  $\varphi$ , etc. As we do not have infinite formulas to express this, let  $\text{E-INT}_G^1(\varphi)$  be an abbreviation for  $\text{E-INT}_G(\varphi)$ , and let  $\text{E-INT}_G^{k+1}(\varphi)$  for  $k \geq 1$  be an abbreviation of  $\text{E-INT}_G(\text{E-INT}_G^k(\varphi))$ . Thus we have  $\text{M-INT}_G(\varphi)$  iff  $\text{E-INT}_G^k(\varphi)$  for all  $k \geq 1$ . It can be shown that the following axiom and rule can be added to the union of  $KD_n$  (the axiom system for individual intentions given in subsection 2.3) and **M1**:

$$\mathbf{M2} \quad \text{M-INT}_G(\varphi) \leftrightarrow \text{E-INT}_G(\varphi \wedge \text{M-INT}_G(\varphi))$$

**RM1** From  $\varphi \rightarrow \text{E-INT}_G(\psi \wedge \varphi)$  infer  $\varphi \rightarrow \text{M-INT}_G(\psi)$  (Induction Rule)

The resulting system is called  $KD_n^{\text{M-INT}_G}$ . Finally, the collective intention is defined by the following axiom:

**M3**  $\text{C-INT}_G(\varphi) \leftrightarrow \text{M-INT}_G(\varphi) \wedge \text{C-BEL}_G(\text{M-INT}_G(\varphi))$

Note that this definition is different from the one given in [30, 11]. The definition would be even stronger if common knowledge instead of common belief appeared in **M3**. However, because common knowledge is almost impossible to establish in multiagent systems due to (among other circumstances) the unreliability of communication media, we do not pursue this strengthening further.

The resulting system, which we call  $KD_n^{\text{C-INT}_G}$ , is the union of the systems  $KD_n^{\text{M-INT}_G}$  (for mutual intentions),  $KD45_n^{\text{C}}$  (for collective beliefs) and axiom **M3**. In [1] we have also defined an alternative version of collective intention that is applicable in different circumstances.

Let us give an informal **example** of the establishment of a collective intention. Two violinists,  $a$  and  $b$ , have studied together and have toyed with the idea of giving a concert together someday. Later this becomes more concrete: they both intend to perform together the solo parts of the Bach Double Concerto ( $\text{INT}(a, \varphi)$  and  $\text{INT}(b, \varphi)$ , where  $\varphi$  stands for “ $a$  and  $b$  perform the solo parts of the Bach Double Concerto”). After communicating with each other about this, they start practicing together. Clearly, a mutual intention as defined in **M2** is now in place (involving nested intentions like  $\text{INT}(a, \text{INT}(b, \text{INT}(a, \varphi)))$  and so on). The communication established a common belief  $\text{C-BEL}_G(\varphi)$  (with  $G = \{a, b\}$ ) about their mutual intention, according to **M3**. As sometimes happens in life, when people are ready, an opportunity appears: Carnegie Hall plans a concert for Christmas Eve, including the Bach Double Concerto. Now they refine their collective intention to a more concrete  $\text{C-INT}_G(\psi)$  (where  $\psi$  stands for “ $a$  and  $b$  perform the solo parts of the Bach Double Concerto at the Christmas Eve concert in Carnegie Hall”). It happens that our two violinists are chosen from among a list of candidates to be the soloists, and both sign the appropriate contract. Because they do this together, we can speak about common knowledge, not merely common belief, of their mutual intention:  $\text{M-INT}_G(\psi) \wedge \text{C-KNOW}_G(\text{M-INT}_G(\psi))$ .

One important difference between common knowledge and common belief is that common knowledge can be justified if needed, and a commonly signed contract provides a perfect basis for this. It is clear that the two violinists have developed a very strong and concrete variant of collective intention due to their common knowledge of the mutual intention.

Even though  $C\text{-INT}_G(\varphi)$  seems to be an infinite concept, collective intentions may be established in practice in a finite number of steps. As defined by **M3**, collective intentions are appropriate to model those situations in which communication, in particular announcements, work, especially if one initiator establishes the team. We have showed in detail in [21, 15] how team formation in such an ideal case may actually work in terms of the first two stages of collective problem solving, namely potential recognition and team formation, and how at these stages the proper attitudes are established through dialogues consisting of the appropriate speech acts.

The above definition is applicable also to cases that may be anticipated when designing a team or system behaviour. For example, emergency situations are such a class of cases. Given a specific application (e.g. a yacht on the sea), different emergency situations are often classified, and based on it, roles of team members are predefined, accordingly. In other words, in specific circumstances, team members know their roles in advance, and have individual intentions to fulfill them. They intend others to fulfill their intentions as well, etc. Thus, the mutual intention  $M\text{-INT}_G(\varphi)$  is in place immediately, especially when saving lives depends on it!

The amount of necessary communication, needed to establish the common belief  $C\text{-BEL}_G(M\text{-INT}_G(\varphi))$ , clearly depends on circumstances, and varies from just recognizing the situation by perception when communication is difficult or impossible, through simply confirming what situation we deal with (then agents' roles are clear), to the more complex cases when, for example, some agents / roles are missing, so that more communication is needed.

The following lemma follows immediately from the definition of collective intention, using lemma 1. It says that a group of agents, when it has a collective intention, is always aware of this.

**Lemma 2.** *Let  $\varphi$  be a formula and  $G \subseteq \{1, \dots, n\}$ . Then the following holds:*

$$C\text{-INT}_G(\varphi) \rightarrow C\text{-BEL}_G(C\text{-INT}_G(\varphi)).$$

One question is if our definition of collective intention does not cover inappropriate cases where real teamwork is out of the question. Bratman [26, Ch. 5] characterizes shared cooperative activity. He presents some examples of situations where the agents share some attitudes, but shared cooperative activity is out of the question. It turns out that our definition excludes these cases, as well. For example,

“Suppose that you and I each intend that we go to New York together, and this is known to both of us. However, I intend that we go together as a result of my kidnapping you and forcing you to join me. The expression of my intention, we might say, is the Mafia sense of “We’re going to New York together”. While I intend that we go to New York together, my intentions are clearly not cooperative in spirit.”

Thus, taking  $\varphi =$ “ $a$  and  $b$  go to New York” with  $a$  for “you”  $b$  for “me”,  $G$  for  $\{a, b\}$ , in the situation above  $E\text{-INT}_G(\varphi)$  and possibly also  $C\text{-BEL}_G(E\text{-INT}_G(\varphi))$  holds, but  $M\text{-INT}_G(\varphi)$  and  $C\text{-INT}_G(\varphi)$  do not. Specifically, it seems unlikely that  $\text{INT}(b, \text{INT}(a, \varphi))$  holds for the Mafioso. Note that Rao, Georgeff and Sonenberg’s *joint intention* among  $G$  to achieve  $\varphi$  is defined as  $E\text{-INT}_G(\varphi) \wedge C\text{-BEL}_G(E\text{-INT}_G(\varphi))$  (translated to our notation), thus it erroneously ascribes a joint intention to go to New York among the agents in the example. Incidentally, a similar one-level definition of mutual goals was also given in [29].

#### 4 A tuning machine for collective commitment

After a group is constituted on the basis of collective intention, another stage of CPS, namely *plan formation* is started, leading ultimately to a *collective commitment* between the team members. While a collective intention may be viewed as an inspiration for team activity, the collective commitment reflects the concrete manner of achieving the intended goal by the team. This concrete manner is provided by planning, and hinges on the allocation of actions according to an adopted plan. This allocation is concluded by agents accepting pairwise (i.e. social) commitments to realize their individual actions. This way, our approach to collective commitments is plan-based.

Let us turn for a moment to the commonsense meaning of teamwork. It is clear that there are different *gradations* of being a team. Take, as **Example 1**, teamwork in a group of researchers who jointly plan their research and divide roles, and who reciprocally keep a check on how the others are doing and help their colleagues when needed in furtherance of their collective intention to prove a theorem. All aspects of teamwork are openly discussed in the team, and members keep each other informed about relevant changes in the plan. Contrast this kind of non-hierarchical teamwork with **Example 2**: a group of spies who all work for the same goal, say to locate Mr. X. In their case a plan is designed by one mastermind, who divides the roles and divulges to each participant *only* the



information that is absolutely necessary for him to do his own part. Thus, members may not know what the main goal is, nor even which other agents are included in the group. In the latter example, even though the connection between members is much looser than in the first one, we would still like to speak about CPS, albeit a non-typical case.

In the two examples above, individual and collective awareness about the ingredients of CPS (like the main goal and the plan to achieve it) ranges from very high in the first example to very low in the second. Very informally, collective commitment is the motivational group attitude that provides the glue needed in a team in order to lead it from a still rather abstract collective intention to concrete team action, and it includes the team members' beliefs about each other and the plan they will follow.

Thus, we claim that the two examples above cannot be covered by *one* generic type of collective commitment. Thus far in the MAS literature, when collective attitudes such as collective or joint intentions and collective or joint commitments were characterized, authors provided just one definition geared towards a typical, ideal type of teamwork [23, 30, 28, 31]. These definitions of collective attitudes were independent of organisational structures and communication possibilities. In contrast, here we will provide a full range of types of collective commitments and weaker group attitudes that play a similar cohesive role, covering the range from proper teams to more loosely connected groups involved in CPS. We also claim that it is important for system developers to make appropriate decisions about the type or gradation of teamwork needed for a given goal in given circumstances, and to have a mechanism that helps them to choose the corresponding type of group commitment to be created.

While investigating the calibration of group commitment to the particular purpose and the specific circumstances, we isolated and separately characterized invariant ingredients of collective commitments. These are:

- collective intention on which the team is built,
- degrees of belief in a team,
- different aspects of team awareness.

They may be viewed as three types of 'dials' that are separately tuned in order to obtain a situation-sensitive notion of collective commitment of a desired strength. Before treating these 'dials', we give a general schema for defining collective commitments. This generic schema together with a tuning mechanism may be viewed as a sort of *tuning machine* for creating collective commitments.

#### 4.1 General schema of collective commitment

In our generic description we will solely define basic ingredients constituting collective commitments, leaving room for case-specific extensions. The obligatory ingredients are related to different aspects of teamwork:

1. Mutual intention  $M\text{-INT}_G(\varphi)$  between a group of agents.  
Let us stress the crucial role of mutual intention when creating a group: the team is *based* on this attitude, and exists as long as the mutual intention exists. Thus, no teamwork is considered without a mutual intention among team members.
2. Social plan  $P$  on which a collective commitment will be based.  
The social plan provides a concrete manner for the team to collectively achieve the overall goal of the system. The predicate  $cons(\varphi, P)$  informally stands for “ $P$  is a correct social plan to achieve  $\varphi$ ”. For a definition and examples of social plans, see 2; for a formal definition of  $cons(\varphi, P)$ , see [8].
3. Pairwise social commitments  $COMM(i, j, \alpha)$  for actions from the plan.  
The group splits the tasks according to their social plan  $P$ , and each agent takes on responsibility to do its part by accepting relevant social commitments.

Next to the above ingredients, different degrees of awareness about them may be present in a team. It may vary from the lack of any awareness to common belief about the given aspect, as it was discussed before. Thus, a general schema covering different types of collective commitment is the following, where the conjuncts between curly brackets may be present or not, according to the position of the awareness ‘dial’ :

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

In words, group  $G$  has a collective commitment to achieve overall goal  $\varphi$  based on social plan  $P$  ( $C\text{-COMM}_{G,P}(\varphi)$ ) iff all of the following hold. The group mutually intends  $\varphi$  (with or without being aware); moreover, successful execution of social plan  $P$  leads to  $\varphi$  ( $cons(\varphi, P)$ ) (with or without the group being aware of this); and finally, for every one of the actions  $\alpha$  from a plan  $P$ , there should be one agent in the group who is socially committed to at least one (mostly other) agent in the group to

fulfil the action ( $\text{COMM}(i, j, \alpha)$ ) (with or without the group being aware of this).

Instantiating the above schema corresponds to tuning the *awareness<sub>G</sub>*-dials from  $\emptyset$ , through individual beliefs and different degrees of  $\text{E-BEL}_G^k$ , to common belief, and/or analogously for degrees of knowledge. These degrees have been discussed in subsection 2.2.

## 4.2 Different aspects of agents' awareness

The notion of collective commitment, whatever strength of it is considered, combines essentially different aspects of teamwork: strictly technical ones related to social plans, as well as those related to agents' intentional stance. The latter concerns different aspects of awareness that appear in a group of agents in the course of CPS. The degree of this awareness, characterized in terms of different types of beliefs, may be different. In the sequel, the strongest version is considered, namely common belief about considered aspect of CPS. For this reason it is justified to speak about *collective awareness* in this context. In other circumstances, the degree of awareness can be weakened by using  $\text{E-BEL}_G$  (or another  $\text{E-BEL}_G^k$ ) instead of  $\text{C-BEL}_G$ . Let us discuss the relevant aspects in detail.

1. Collective intention is the attitude constituting the group as a whole. Thus, it introduces (rather strong) collective awareness of the group as a cooperative team of agents. Formally this is expressed as a conjunct in the definition of collective intention:

$$\text{C-BEL}_G(\text{M-INT}_G(\varphi))$$

2. When a team of agents exists, the next step is plan generation or adoption. Regardless of the method of arriving at this point, the type of awareness connected with this is collective awareness of the correctness of the plan with respect to the overall goal. Formally:

$$\text{C-BEL}_G(\text{cons}(\varphi, P))$$

3. When a plan as a recipe is in place, then the particular actions from it need to be allocated to particular team members in order to create pairwise social commitments between them. This way a social structure is built within a team, and the plan acquires the property of being social. The type of awareness connected with this phase may be twofold.

- (a) The first one is a collective awareness of the social structure in a team with respect to a given plan. This includes a *detailed* collective awareness of each social commitment involved. Formally:

$$\bigwedge_{\alpha \in P} \bigvee_{i, j \in G} \text{C-BEL}_G(\text{COMM}(i, j, \alpha))$$

This corresponds to the interpretation *de re*.

- (b) The second one refers to a more *global* collective awareness of the social structure within the team, namely of the bare existence of social commitments with respect to a given social plan. Formally:  

$$\text{C-BEL}_G(\bigwedge_{\alpha \in P} \bigvee_{i,j \in G} \text{COMM}(i, j, \alpha))$$
This corresponds to the interpretation *de dicto*.

The distinction *de re* / *de dicto* stems from the philosophy of language [32]. A sentence of the form  $\exists x \text{BEL}(j, A(x))$  is a *de re* belief attribution which relates  $j$  to a *res*, an individual that the belief is about. On the other hand,  $\text{BEL}(j, \exists x A(x))$  is a *de dictum* belief attribution, relating  $j$  to a *dictum*, namely the proposition  $\exists x A(x)$ . This distinction is also fruitful for complex epistemic operators such as common belief. Note that  $\text{C-BEL}_G$  in (a) and (b) distributes over conjunction ( $\bigwedge_{\alpha \in P}$ ), so that only the position of  $\text{C-BEL}_G$  with respect to  $\bigvee_{i,j \in G}$  matters.

The above aspects of awareness will be viewed as building blocks when distinguishing different strengths of collective commitments.

### 4.3 Different notions of collective commitment

The following exemplar definitions are produced by keeping the *awareness<sub>G</sub>*-dial fixed to a choice between  $\emptyset$  and collective belief, and the dial for ‘kind of mutual intention’ fixed as the standard definition of subsection 3.1. We will start from the strongest form of collective commitment, fully reflecting the collective aspects of CPS. Subsequently, some underlying assumptions will be relaxed, leading ultimately to weaker notions of team and distributed commitment.

**Robust collective commitment** Our discussion on different types on collective commitments will start from the two cases based on collective planning. Additionally, for every one of the actions  $\alpha$  that occur in social plan  $P$ , there should be one agent in the group who is socially committed to at least one (mostly other) agent in the group to fulfil the action. Moreover, in a *robust collective commitment* ( $\text{R-COMM}_{G,P}$ ), there is a detailed awareness of social commitments in the team:

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

By the last conjunct, everybody’s responsibility is public. The aspect of sharing responsibility is of crucial importance here. Among others it

implies that there is no need for an initiator in such a team.

**Example** Robust collective commitment may be applicable in (small) companies where all team members involved are share-holders. Typically, planning is done collectively, whether from first principles or choosing from a plan library. Everybody's responsibility is public, because the social commitments are established publicly. In particular, when any form of revision is needed due to dynamic circumstances, the entire team may be collectively involved.

**Strong collective commitment** In contrast to robust collective commitment, in the case of *strong collective commitment* (S-COMM<sub>G,P</sub>), there is a global awareness about particular social commitments, but the group as a whole believes that things are under control, i.e., that every part of the plan is within somebody's responsibility:

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

As the responsibility is not shared due to the lack of detailed awareness in the last conjunct, the case of a team leader or initiator fits here. Also, as pairwise social commitments are not collectively known, they cannot be collectively revised when such a need appears.

**Example** Strong collective commitment may be applicable in companies with one or more leaders and rather separate subteams. Typically, planning is done collectively. However, establishing bilateral commitments is not done publicly in the whole team, but in subgroups. Sometimes this suffices, and it is sometimes preferable in order not to waste energy.

**Weak collective commitment** In the weaker cases of collective commitment, the degree of team awareness is even more limited. When the plan as a whole is not known to the team, for example, if no collective decision making is assumed, there is no awareness that the plan leads to proper realization of the goal (C-BEL<sub>G</sub>cons(φ, P) is not in place). In this case we deal with a *weak collective commitment* (W-COMM<sub>G,P</sub>):

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

In this case, the team knows the overall goal, but does not know details of the plan: there is no collective awareness of the plan's correctness. Apparently, also in this case no collective revision of social commitments may take place.

**Example** Weak collective commitment may be applicable in companies with a dedicated planner or planning department. Typically, the planner individually does believe in the plan's correctness  $cons(\varphi, P)$ , and this may suffice.

**Team commitment** In the case of *team commitment* (T-COMM $_{G,P}$ ) agents remain aware solely about their piece of work, without any orientation about involvement of others. In this situation, there is no common belief that all actions have been adopted by other committed members, but a team as a structure still exists :

$$\text{T-COMM}_{G,P}(\varphi) \leftrightarrow \text{C-INT}_G(\varphi) \wedge cons(\varphi, P) \wedge \bigwedge_{\alpha \in P} \bigvee_{i,j \in G} \text{COMM}(i, j, \alpha)$$

Because of the presence of collective intention, the overall goal and composition of the team are collectively believed. Planning is not at all collective: it may be that even task division is not public; this is often done on purpose. Thus, distribution of social commitments cannot be public either.

**Example** Team commitment may be applicable in companies assigning limited trust to its employees. Information about the precise involvement of colleagues and other aspects of the plan may be confidential.

**Distributed commitment** The last case distinguished here is *distributed commitment* (D-COMM $_{G,P}$ ). It deals with the situation when agents' awareness is even more restricted: they may not even know the overall goal, only their share in an 'undefined' project:

$$\text{D-COMM}_{G,P}(\varphi) \leftrightarrow cons(\varphi, P) \wedge \bigwedge_{\alpha \in P} \bigvee_{i,j \in G} \text{COMM}(i, j, \alpha).$$

This means that no 'real' team of cooperating agents is created, so that no collective intention C-INT $_G$  is in place. Instead, a rather loosely coupled group of agents works in a distributed manner without autonomous involvement in the project to be realized.

**Example** Distributed commitment may be applicable in companies contracting out some labour to outsiders. The overall goal and the group

of agents involved may be classified information, e.g. in order to avoid competition.

## 5 Discussion and conclusions

This article falls within a larger research program, the first part of which presents a *static* characterization of CPS with collective commitment as a central notion [1]. This theory is built incrementally starting from individual intentions, through social commitments, leading ultimately to the collective level of intentions and commitments. All these notions play a crucial role in practical reasoning. They are defined in multi-modal logics with clear semantics (cf. [1]), comprising a descriptive view on collective motivational attitudes.

In this paper, we give necessary and sufficient conditions for such collective motivational attitudes to be present. In this way, we hope to make the behavior of a team easier to predict. We have tried to find *minimal* conditions for collective intentions to be present, and not to weigh down the definitions with all aspects that play a part in the establishment of collective intentions. Such elements as conventions, abilities, opportunities, power relations and social structure (see [33, 27, 23] for a thorough discussion) certainly are important, and we leave open the possibility of defining and using them in specific cases where they play a crucial role. For example, abilities and opportunities are important in the dialogues leading up to the establishment of a collective intention and a team based on it [15]. Power relations and social structure, on the other hand, are reflected in the definitions of collective commitments.

The definitions of collective commitments are not overloaded, and therefore easy to understand and to use. In contrast to [30], we do not give one iron-clad definition of collective commitment here. Instead, we provide a sort of tuning mechanism for the system developer to calibrate an appropriate type of collective commitment, taking into account both the circumstances in which a group is acting, for example possibilities of communication, as well as organizational structure. Such an approach is especially strong when re-planning is needed. The multi-modal logic framework allows to express subtle aspects of CPS, modeling different situations occurring in practical domains.

The presented system defining collective intentions and collective commitments is decidable, however it is highly complex. In order to decide whether a given formula is valid, one needs an amount of time that grows exponentially in terms of the length of the formula in question [34]. Thus,

in general it is not feasible to give automated proofs of desired properties; at least there is no single algorithm that performs well on all inputs. As with other modal logics, the better option would be to develop a variety of different algorithms and heuristics, each performing well on a limited class of inputs. Also, when considering specific applications it is possible to reduce some of the infinitary character of collective beliefs and intentions to more manageable proportions (cf. [13, Ch. 11]).

The static definitions of relevant motivational attitudes given express solely vital aspects of CPS, leaving room for case-specific extensions. This set of *teamwork axioms* constitutes a definition of motivational attitudes in BDI systems. A specification of their evolution in a dynamic environment has been presented in [8]. This way they may serve a system designer as a specification to create a correct and complete system, as well as to verify a system.

The presented analysis of social and collective attitudes in teams of agents assumes a rather high level of idealization: solely strictly cooperative teams are considered. This leads to a strong definition of collective intention, based on agents' mutual intentions, and then to a plan-based collective commitment. Even though in [35, 2] we introduced a general tuning mechanism, presented here in section 4, to calibrate the strength of collective commitments fitting to a variety of circumstances, an essential ingredient of these definitions — agents' awareness — is formalized by means of a strong notion of common belief. We agree with Cristiano Castelfranchi that after investigating and formalizing this basic case, it is time to relax some of the strong assumptions underlying this research in order to take a closer look on weaker and more distributed forms of cooperation.

Let us finally mention an important cognitive aspect of our current research. According to Herbert Simon, who coined the term *bounded rationality*, “boundedly rational agents experience limits in formulating and solving complex problems and in processing (receiving, storing, retrieving, transmitting) information” [36]. We agree with Simon that models which present humans as logically omniscient or as perfectly rational in the sense of optimizing their own utility are problematic. Moreover, we would propose to extend this discussion to software agents, because they need to reason under bounded rationality as well as humans.

In [37], we take the point of view that for successful coordination and cooperation in a multiagent system, participants need to establish a sufficiently accurate *awareness* of the current situation. Awareness is understood here as a limited form of consciousness: in the minimal form, it



refers to the state of an agent's beliefs about itself, about others and about the environment. When considered in the context of agents' mental states, this leads to distinguishing three levels of awareness: intra-personal (about the agent itself), inter-personal (about other agents as individuals), and group awareness.

Problems in modeling agents' awareness on all three levels are analyzed in [37]. It turns out that both the communication medium and agents' cognitive and computational limitations make the achievement of awareness difficult. Cognitive science is used to analyze and explain problems in human awareness, based on the concept of bounded rationality. The BDI framework turns out to be fruitful both for precisely formulating the problems concerning agents' awareness, and, even more importantly, for formulating and comparing possible solutions. Thus, the two fields of cognitive science and MAS mutually benefit from each other's viewpoints, especially in the light of the currently growing need for teamwork in which both computational agents and humans are involved.

## References

1. Dunin-Kęplicz, B., Verbrugge, R.: Collective intentions. *Fundamenta Informaticae* **51(3)** (2002) 271–295
2. Dunin-Kęplicz, B., Verbrugge, R.: A tuning machine for cooperative problem solving. *Fundamenta Informaticae* **63** (2004) 283–307
3. Bratman, M.: *Intention, Plans, and Practical Reason*. Harvard University Press, Cambridge (MA) (1987)
4. Weiss, G., ed.: *Multiagent Systems*. MIT Press, Cambridge, MA (1999)
5. Levesque, H., Cohen, P., Nunes, J.: On acting together. In: *Proceedings Eighth National Conference on AI (AAAI90)*, Menlo Park (CA), Cambridge (MA), AAAI-Press and MIT Press (1990) 94–99
6. Cohen, P., Levesque, H.: Intention is choice with commitment. *Artificial Intelligence* **42** (1990) 213–261
7. Dunin-Kęplicz, B., Verbrugge, R.: A reconfiguration algorithm for distributed problem solving. *Engineering Simulation* **18** (2001) 227 – 246
8. Dunin-Kęplicz, B., Verbrugge, R.: Evolution of collective commitments during teamwork. *Fundamenta Informaticae* **56** (2003) 329–371
9. Ferber, J.: *Multi-agent Systems: An Introduction to Distributed Artificial Intelligence*. Addison Wesley, Reading (MA) (1999)
10. Rao, A., Georgeff, M.: Modeling rational agents within a BDI-architecture. In Fikes, R., Sandewall, E., eds.: *Proceedings of the Second Conference on Knowledge Representation and Reasoning*, Morgan Kaufman (1991) 473–484
11. Dunin-Kęplicz, B., Verbrugge, R.: Collective motivational attitudes in cooperative problem solving. In Gorodetsky, V., ed.: *Proceedings of the First International Workshop of Eastern and Central Europe on Multi-agent Systems (CEEMAS'99)*, St. Petersburg (1999) 22–41
12. Blackburn, P., de Rijke, M., Venema, Y.: *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge (2001)

13. Fagin, R., Halpern, J., Moses, Y., Vardi, M.: Reasoning about Knowledge. MIT Press, Cambridge, MA (1995)
14. Meyer, J.J.C., van der Hoek, W.: Epistemic Logic for AI and Theoretical Computer Science. Cambridge University Press, Cambridge (1995)
15. Dignum, F., Dunin-Kępicz, B., Verbrugge, R.: Creating collective intention through dialogue. *Logic Journal of the IGPL* **9** (2001) 145–158
16. Dignum, F., Conte, R.: Intentional agents and goal formation: Extended abstract. In Singh, M., Rao, A., Wooldridge, M., eds.: Proceedings Fourth International Workshop on Agent Theories, Architectures and Languages. Volume 1365 of Lecture Notes in Computer Science., Berlin, Springer Verlag (1997) 231–243
17. Holyoak, K.J., Morrison, R.G.: The Cambridge Handbook of Thinking and Reasoning. Cambridge University Press, Cambridge (2005)
18. Castelfranchi, C.: Commitments: From individual intentions to groups and organizations. In Lesser, V., ed.: Proceedings First International Conference on Multi-Agent Systems, San Francisco, AAAI-Press and MIT Press (1995) 41–48
19. Segerberg, K.: Bringing it about. *Journal of Philosophical Logic* **18** (1989) 327–347
20. Castelfranchi, C., Miceli, M., Cesta, A.: Dependence relations among autonomous agents. [38]
21. Dignum, F., Dunin-Kępicz, B., Verbrugge, R.: Agent theory for team formation by dialogue. In Castelfranchi, C., Lesperance, Y., eds.: Intelligent Agents VII: Agent Theories, Architectures and Languages. Volume 1986 of Lecture Notes in Computer Science., Berlin, Springer Verlag (2001) 141–156
22. Jennings, N.: Commitments and conventions: The foundation of coordination in multi-agent systems. *Knowledge Engineering Review* **3** (1993) 223–250
23. Wooldridge, M., Jennings, N.: Cooperative problem solving. *Journal of Logic and Computation* **9** (1999) 563–592
24. Haddadi, A.: Communication and Cooperation in Agent Systems: A Pragmatic Theory. Volume 1056 of LNAI. Springer Verlag, Berlin (1995)
25. Tuomela, R., Miller, K.: We-intentions. *Philosophical Studies* **53** (1988) 367–390
26. Bratman, M.: Faces of Intention. Cambridge University Press, Cambridge (1999)
27. Tuomela, R.: The Importance of Us: A Philosophical Study of Basic Social Notions. Stanford Series in Philosophy. Stanford University Press, Stanford (CA) (1995)
28. Rao, A., Georgeff, M., Sonenberg, E.: Social plans: A preliminary report. [38] 57–76
29. Wooldridge, M., Jennings, N.: Towards a theory of collective problem solving. In Perram, J., Muller, J., eds.: Distributed Software Agents and Applications. Volume 1069 of LNAI. Springer Verlag, Berlin (1996) 40–53
30. Dunin-Kępicz, B., Verbrugge, R.: Collective commitments. In Tokoro, M., ed.: Proceedings Second International Conference on Multi-Agent Systems, Menlo Park (CA), AAAI-Press (1996) 56–63
31. Grosz, B., Kraus, S.: The evolution of SharedPlans. In Rao, A., Wooldridge, M., eds.: Foundations of Rational Agency. Kluwer, Dordrecht (1999) 227–262
32. Quine, W.: Quantifiers and propositional attitudes. *Journal of Philosophy* **53** (1956) 177–187
33. Singh, M.: Commitments among autonomous agents in information-rich environments. In Boman, M., de Velde, W.V., eds.: Multi-Agent Rationality (Proceedings of MAAMAW'97). Volume 1237 of LNAI. Springer Verlag, Berlin (1997) 141–155
34. Dziubinski, M., Verbrugge, R., Dunin-Kępicz, B.: Complexity of a theory of collective attitudes in teamwork. In Skowron, A., Barthes, J.P., Jain, L., Sun, R., Morizet-Mahoudeaux, P., Liu, J., Zhong, N., eds.: Proceedings of the 2005

- IEEE/WIC/ACM International Conference on Intelligent Agent Technology, Los Alamitos (CA), IEEE Computer Society (2005) 579–586
35. Dunin-Kępicz, B., Verbrugge, R.: Calibrating collective commitments. In Marik, V., Mueller, J., Pechoucek, M., eds.: Proceedings of the 3rd International Central and Eastern European Conference on Multi-Agent Systems. Volume 2691 of LNAI., Berlin, Springer Verlag (2003) 73–83
  36. Williamson, O.: The economies of organization: The transaction cost approach. *American Journal of Sociology* **87** (1981) 548–577
  37. Dunin-Kępicz, B., Verbrugge, R.: Awareness as a vital ingredient of teamwork. In Weiss, G., Stone, P., eds.: Proceedings of the Fifth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06), Los Alamitos, IEEE / ACM Press (2006)
  38. Werner, E., Demazeau, Y., eds.: Decentralized A.I.-3. In Werner, E., Demazeau, Y., eds.: Decentralized A.I.-3, Amsterdam, Elsevier (1992)