Case-study for TeamLog, a theory of teamwork

Barbara Dunin-Kęplicz, Rineke Verbrugge, and Michał Ślizak

Abstract This article presents a case-study of a theoretical multi-agent system designed to clean up ecological disasters. It focuses on the interactions within a team of agents, outlines their goals, and establishes the necessary distribution of knowledge and commitment throughout the team. In addition, cooperation between teams is treated. These aspects of teamwork are presented in the TEAMLOG formalism. We will show how to make a bridge between theoretical foundations of a BDI system and a real application. Complex team attitudes are justified to be necessary in the course of teamwork. At the same time we show how to establish them on a sufficient, but still minimal level.

1 Defining teamwork

When constructing BDI systems, firstly a model of an agent as an *individual, autonomous* entity [21] has to be constructed. Nowadays a key point is to organize agents' cooperation in a way allowing the achievement of their common goal, while preserving their individual autonomy (see [18, 23, 15, 24, 1, 6, 7, 14] for some logical approaches to teamwork). The BDI model comprises agents' individual beliefs, goals, and intentions. However in teamwork, when a team of agents needs to work together in a planned and coherent way, agents' individual attitudes are not enough: the group needs to present a common attitude over and above individual ones. This group attitude is a necessary condition for a loosely-coupled group of agents to become a strictly cooperative team. In this case-study we focus on full cooperation, where agents' attitudes are considered on the individual, social (i.e. bilateral) and collective level.

Barbara Dunin-Kęplicz
Institute of Informatics, Warsaw University, Banacha 2, 02-097 Warsaw, Poland
Institute of Computer Science, Polish Academy of Sciences, Ordona 21, 1-237 Warsaw, Poland
e-mail: keplicz@mimuw.edu.pl

Michał Ślizak
Institute of Informatics, Warsaw University, Banacha 2, 02-097 Warsaw, Poland, e-mail: Michal.Slizak@mimuw.edu.pl

Rineke Verbrugge
Department of Artificial Intelligence, University of Groningen, P.O. Box 407, 9700 AK Groningen, The Netherlands, e-mail: rineke@ai.rug.nl

A theory of individual and group beliefs has been formalized in terms of epistemic logic [12, 19, 20]. General, common, and distributed knowledge and belief were defined in terms of agents' individual knowledge and belief. Different axiom systems express various properties of knowledge and belief, while the corresponding semantics naturally reflect these properties.

As regards motivational attitudes, the situation is much more complex as the bilateral and collective notions cannot be viewed as a sort of sum of individual ones. Additional subtle and diverse aspects of teamwork need to be isolated and appropriately defined. The *static, descriptive* theory of collective motivational attitudes TEAMLOG [6, 7, 11] has been formed on the basis of individual goals, beliefs and intentions of cooperating agents. It addresses the question what it means for a group of agents to have a *collective intention*, and then a *collective commitment* to achieve a common goal. While collective intention consolidates a group as a strictly cooperating team, collective commitment leads to team action, i.e., to coordinated realization of individual actions by committed agents according to a plan. The social plan can be constructed from first principles, or may be chosen from a repository of pre-constructed plans. Both collective intentions and collective commitments allow to fully express the potential of strictly cooperative teams [6, 7].

When modelling group attitudes, agents' awareness about the overall situation needs to be taken into account. Awareness is understood here as a limited form of consciousness: it refers to the state of an agent's beliefs about *itself* (*intra-personal*), about *others* (*inter-personal*) and about *the environment* (*group awareness*). Thus, various epistemic logics and different gradations of group information (from distributed belief to common knowledge) are adequate to formalize agents' awareness [12, 7, 20].

In TEAMLOG, group awareness is usually expressed in terms of *common belief*, (C-BEL_G) fully reflecting collective aspects of agents' behavior. Due to its infinitary flavor, this concept has a high complexity: its satisfiability problem is EXPTIME-complete [11]. There are general ways to reduce the complexity by restricting the language, by allowing only a small set of atomic propositions or restricting the modal context in formulas, as proved in [11, 10]. However, when building MAS applications, it may be more profitable to use domain-specific means to tailor TEAMLOG to the circumstances in question, calling for weaker forms of awareness [9]. In this case-study of prevention of ecological disasters, we will illustrate how to adjust TEAMLOG to a specific environment. Our aim is to show how the infinitary definitions of collective attitudes can be reduced in a real-world situation.

This paper is structured into several sections. This one introduces the problem in general terms. Next, some definitions and assumptions regarding the environment are presented, including an outline of the interactions within and between teams. This is followed in section 3 by definitions of plans to be executed. In section 4, a short reminder of part of TEAMLOG is given. In section 5 we explore the minimal requirements for successful teamwork in the case-study from the theoretical perspective. A short discussion sums up this case-study.

2 The case-study: ecological disasters

The case-study deals with prevention and repair of ecological disasters caused by specific poisons, by means of heterogeneous multi-agent teams. They work in situations where time is critical and resources are bounded [17, 22]. The maintenance goal *safe* is to keep a given region *REG* safe or to return it to safety if it is in danger.

Possible hazards are two kinds of poison, X_1 and X_2, which are dangerous in high concentrations. They may be explosive if they react with one another to form compound $X_1 \oplus X_2$, which happens at high concentrations. Three functions f_1, f_2 and f_3 reflect the influence of temperature $t(A)$, pressure $p(A)$ and concentrations $c_1(A)$ and $c_2(A)$ of poisons X_1 and X_2 at location A on the possible danger level at that location. The function ranges are divided into three intervals, as follows:

The first poison X_1:

- $safe_1$ iff $f_1(p(A), t(A), c_1(A)) \in [0, v_1]$;
- $risky_1$ iff $f_1(p(A), t(A), c_1(A)) \in (v_1, n_1]$;
- $dangerous_1$ iff $f_1(p(A), t(A), c_1(A)) \in (n_1, \infty)$;

The second poison X_2:

- $safe_2$ iff $f_2(p(A), t(A), c_2(A)) \in [0, v_2]$;
- $risky_2$ iff $f_2(p(A), t(A), c_2(A)) \in (v_2, n_2]$;
- $dangerous_2$ iff $f_2(p(A), t(A), c_2(A)) \in (n_2, \infty)$;

The compound poison $X_1 \oplus X_2$:

- $safe_3$ iff $f_3(p(A), t(A), c_1(A), c_2(A)) \in [0, v_3]$;
- $risky_3$ iff $f_3(p(A), t(A), c_1(A), c_2(A)) \in (v_3, n_3]$;
- *explosive* iff $f_3(p(A), t(A), c_1(A), c_2(A)) \in (n_3, \infty)$;

We define $safe := safe_1 \wedge safe_2 \wedge safe_3$ and refer to it as a goal and as a predicate. There are also thresholds ε_1 and ε_2: when the concentration of a poison X_i exceeds ε_i, the respective function f_i is computed.

2.1 Starting point: the agents

This model reflects cooperation between humans, software agents, robots, unmanned aerial vehicles (*UAVs*) [4, 3], and a helicopter steered by a pilot. The whole process is coordinated by one *coordinator*, who initiates cooperation, coordinates teamwork between different teams, is responsible for dividing the disaster zone into sectors and assigning a team to each sector to perform clean-up. Several teams of similar make-up work in parallel, aiming to prevent or neutralize a contamination. Each of these teams consist of:

- one *UAV* - responsible to the coordinator for keeping assigned sectors in a safe state. Cannot carry heavy load, but has considerable computational capabilities for planning and is capable of mapping terrain and observation;
- one regular helicopter steered by the human *pilot*@@, can independently choose the order in which it will clean up assigned areas@@;
- n identical neutralizing robots $rob_1, \ldots, rob_n$ - responsible to their *UAV* for cleaning up a zone.

2.2 Cooperation between teams

The entire disaster zone is divided into sectors by the coordinator, based on terrain type, team size and known hot spots. Teams are responsible for (possibly many) sectors. Each team's *UAV* prepares a plan to keep its sectors safe. Each plan is judged based on a fitting function *fit*, which must take into account:

1. available robots, current task, load, capacity and position of each one,
2. whether the plan relies on robots from other teams to help the current one,
3. task priorities,
4. the minimum amount of time it takes to implement,
5. the minimum amount of robots it requires.

The *UAVs* communicate and cooperate. If performing tasks requires more robots than are currently available, an *UAV* can call for reinforcements from another *UAV*. Of course fulfilling one's team's objectives has a priority over helping others.

2.3 A bird's-eye view on cases

To maintain the goal *safe*, the situation is monitored on a regular basis with frequency *freq*. During *situation recognition*, in the risky cases monitoring is performed twice as frequently. Depending on the mixture and density of poisons in a location, some general cases followed by the relevant procedures are established. All remedial actions are to be performed relative to the contaminated area:

Case *safe***:**
  true $\longrightarrow$ *situation recognition*
Case *dangerous*$_1$**:**
  rain $\longrightarrow$ *liquid* L_1 *to be poured on the soil*
  normal or dry $\longrightarrow$ *liquid* L_2 *to be sprayed from the air*
Case *dangerous*$_2$**:**
  rain $\longrightarrow$ *solid* S_1 *to be spread*, followed by *liquid catalyst* K_1 *to be poured*
  normal or dry $\longrightarrow$ *solid* S_1 *to be spread*
Case *explosive***:**
  before explosion $\longrightarrow$ *evacuation*
  after explosion $\longrightarrow$ *rescue action*

Due to the space limit, we cannot present too many details of plans. Failure handling will not be discussed for the same reason.

3 Global plans

In order to control the amount of interactions and decrease the time needed to establish beliefs, the accepted team model is hierarchical. A coordinator views a team as a single cleaning robot, even though the *UAVs* use many autonomous neutralizing robots to perform their work.

3.1 The social plan $\langle Cleanup \rangle$

The social plan for which the coordinator and *UAVs* are responsible, is designed with respect to location A. It is a while-loop, in which observation is interleaved with treatment of current dangers by level of priority, from most to least dangerous. The goal (denoted as $Clean$) is to keep locations in a *safe* state.

begin
  $freq := a$; $\{freq$ - interval between two checks of the environment$\}$

while true do
  $\langle$ Plan $SR\rangle$ {Compute the situation at A, with frequency $freq$}
  if *explosive* then do $\langle$ Plan E $\rangle$ end;
  elif $dangerous_1$ and rain then do $\langle$ Plan D_1R $\rangle$ end;
  elif $dangerous_1$ then do $\langle$ Plan D_1N $\rangle$ end;
  elif $dangerous_2$ and rain then do $\langle$ Plan D_2R $\rangle$ end;
  elif $dangerous_2$ then do $\langle$ Plan D_2N $\rangle$ end;
  elif $risky_1 \vee risky_2 \vee risky_3$ then $freq := \frac{a}{2}$ end
  else {safe situation} $freq := a$ end;
end
end.

3.2 The social plan $\langle SR\rangle$

This plan performs situation recognition at location A.

begin
  $C_1 := c_1(A)$ {C_1 is the measured concentration of poison X_1 at A}
  $C_2 := c_2(A)$ {C_2 is the measured concentration of poison X_2 at A}
  $T := t(A)$ {T is the measured temperature at A}
  $P := p(A)$ {P is the measured air pressure at A}
  {Computation of the situation at A}
  if $C_1 > \varepsilon_1$ then compute $f_1(C_1, T, P)$ end;
  if $C_2 > \varepsilon_2$ then compute $f_2(C_2, T, P)$ end;
  if $C_1 > \varepsilon_1$ and $C_2 > \varepsilon_2$ then compute $f_3(C_1, C_2, T, P)$ end;
end.

After an explosion, evacuation and rescue of people @@should take place@@. This subject is discussed in many studies [22, 17] and will not be elaborated here. See the Appendix for the other plans included in $\langle Cleanup\rangle$.

4 TeamLog: a logical theory for teamwork

It does not suffice for agents to only have an individual intention (INT) towards their projection of the social plan. They would still act as individuals, so if something new appears in their region or the circumstances change calling for re-planning, the group would be helpless to adapt, as not being formed properly. Thus, group attitudes such as *collective intentions* and *collective commitments* become relevant. Due to the space limit, only a short reminder of TEAMLOG can be given here. For extensive explanations and discussion, see [6, 7].

4.1 Definitions of beliefs in TeamLog: general and common belief

For the individual part, we adopt a standard $KD45_n$ system for n agents governing individual belief operator BEL, as explained in [12]. Additionally, for group beliefs, with $G \subseteq \{1, \ldots, n\}$, as in [12]:

C1 $\text{E-BEL}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} BEL(i, \varphi)$ *(general belief: "everyone believes")*

And in more general iterated form:

$$\text{E-BEL}^k_G(\varphi) \leftrightarrow \text{E-BEL}^{k-1}_G(\text{E-BEL}_G(\varphi)), \text{ where } \text{E-BEL}^1_G(\varphi) \equiv \text{E-BEL}_G(\varphi)$$

$\text{C-BEL}_G(\varphi)$ - "it is common belief in the group G that φ is true":

C2 $\text{C-BEL}_G(\varphi) \leftrightarrow \text{E-BEL}_G(\varphi \wedge \text{C-BEL}_G(\varphi))$
RC1 From $\varphi \to \text{E-BEL}_G(\psi \wedge \varphi)$ infer $\varphi \to \text{C-BEL}_G(\psi)$ *(induction)*

4.2 Definitions of intentions in TeamLog: collective intention

For operator INT, the TEAMLOG axioms comprise the system KD_n, including the intention consistency axiom D. In addition, the system developer may choose whether or not to add positive and negative introspection of intentions (see [6]).

The focus in TEAMLOG on strictly cooperative teams makes the definition of collective intention rather strong. It is certainly not sufficient that all members of the team G have the associated individual intention $\text{INT}(i, \varphi)$ to achieve φ, i.e. a *general intention*. To exclude 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 this a *mutual intention* ($\text{M-INT}_G(\varphi)$). Furthermore, all team members are aware of this mutual intention by a common belief: $\text{C-BEL}_G(\text{M-INT}_G(\varphi))$. Of course, team members remain autonomous in maintaining their other motivational attitudes, and may compete about other issues.

M1 $\text{E-INT}_G(\varphi) \leftrightarrow \bigwedge_{i \in G} \text{INT}(i, \varphi)$ *(general intention: "everyone intends")*

We also iteratively define E-INT^k_G, similar to higher-order general beliefs:

M1′ $\text{E-INT}^k_G(\varphi) \leftrightarrow \text{E-INT}^{k-1}_G(\text{E-INT}_G(\varphi))$, where $\text{E-INT}^1_G(\varphi) \equiv \text{E-INT}_G(\varphi)$

Mutual and collective intentions are governed by:

M2 $\text{M-INT}_G(\varphi) \leftrightarrow \text{E-INT}_G(\varphi \wedge \text{M-INT}_G(\varphi))$ *(mutual intention)*
RM1 From $\varphi \to \text{E-INT}_G(\psi \wedge \varphi)$ infer $\varphi \to \text{M-INT}_G(\psi)$ *(induction)*
M3 $\text{C-INT}_G(\varphi) \leftrightarrow \text{M-INT}_G(\varphi) \wedge \text{C-BEL}_G(\text{M-INT}_G(\varphi))$ *(collective intention)*

4.3 Awareness and tuning collective attitudes

The definitions of collective attitudes from TEAMLOG allow to calibrate the strength of agents' *awareness*. In this case-study we adapt the strength of collective attitudes to the specific domain, and to set relevant group attitudes at the minimal level ensuring the effective team operation. Here the notion of awareness [9] comes in very useful. In general, the question regarding the level of awareness about each specific aspect of teamwork needs to be addressed.

@@The level of awareness@@ is defined by the strength of agents' beliefs. Instances of $awareness_G$ in definitions of motivational attitudes can be anything from $\emptyset$, through individual beliefs, different levels of E-BEL^k_G, to common belief C-BEL_G. Stronger levels of belief may increase communication, since beliefs have to be propagated. It has been argued that in ideal teamwork, $awareness_G$ is taken to be C-BEL_G [6]. Supposing that the communication medium is perfect, it is possible to attain this. In practical implementations (in an asynchronous, uncertain medium) common knowledge (C-KNOW_G) has been proven to be impossible to achieve [16], and common belief (C-BEL_G) under extremely restricted constraints [8]; usually only a finite approximation E-BEL^k_G can be achieved. Here follows the flexible scheme for collective intentions:

M3′ $\text{C-INT}_G(\varphi) \leftrightarrow \text{M-INT}_G(\varphi) \wedge \mathit{awareness}_G(\text{M-INT}_G(\varphi))$

4.4 Collective commitment in TeamLog

After a group is constituted on the basis of collective intention, a *collective commitment* between the team members needs to be established. While a collective intention is an inspiration for team activity, the collective commitment reflects the concrete manner of achieving the goal. This concrete manner is provided by planning, and hinges on the allocation of actions according to a social plan. This allocation is concluded in bilateral (i.e. social) commitments to realize their individual actions. This way, our approach to collective commitments is plan-based. A bilateral commitment from agent i towards agent j to perform action α is represented as $\text{COMM}(i, j, \alpha)$; in this bilateral commitment is viewed as a primitive notion, but see [7] for its characterization and governing axiom.

Collective commitment schema

A flexible tuning schema for *collective commitments* is presented in [7]. In words, group G has a *collective commitment* to achieve goal φ based on social plan P ($\text{C-COMM}_{G,P}(\varphi)$) iff all of the following hold (in the corresponding definition below, parts between curly brackets may or may not be present): The group mutually intends φ (with or without being aware); moreover, successful execution of social plan P leads to φ ($\mathit{cons}(\varphi, P)$) (with or without the group being aware of this); and finally, for every one of the actions α from a plan P, there should be one agent in the group who is bilaterally committed to another agent in the group to fulfil the action ($\text{COMM}(i, j, \alpha)$) (with or without the group being aware of this):

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

Weak collective commitment

Different types of collective commitments related to different organization structures and environments have been introduced in terms of a ‘tuning machine’ [7]. One instantiation of the above scheme by tuning the ‘awareness dial’ is the *weak collective commitment* [7]. 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, so no $\text{C-BEL}_G(\mathit{cons}(\varphi, P)))$, even though there is a global awareness that things are under control. Weak collective commitment may be applicable in teams with a dedicated planner, who takes care of the plan’s correctness $\mathit{cons}(\varphi, P)$:

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

5 Adjusting the TeamLog definitions to the case-study

Why is a collective intention and a social plan still not enough to start team action in the case-study? Because agents may not feel *responsible* for their share. Thus, they need to commit to performing their part from the plan. Here, agents naturally commit to the coordinator, either directly or via an *UAV* acting as a ‘middle manager’.

Now, what is the type of collective commitment fitting to the scenario? On the inner-team level, the *UAV* has the highest level of awareness within its team as it knows the entire social plan. There is no need for the others to know all the details. This corresponds to *weak collective commitment*, as defined above.

On the inter-team level, the controller has the highest awareness. The *UAVs* he oversees only need to know their part of the overall plan, and believe that it has been shared among all *UAVs*. The controller knows both the plan and task distribution.

5.1 Organization structure: who is socially committed to whom?

The coordinator is socially committed to achieving the main goal, with respect to the social plan as a whole. Other agents are committed to their share in the plan. The coordinator is committed towards itself and towards the relevant control authority, e.g. the national environmental agency for which he works.

The *UAV* is committed towards the coordinator with respect to achieving his part of the plan - keeping specified regions in a safe state.

The robots commit to perform their share to their leading *UAV*, which has the power to uncommit them. There is a clear hierarchy where the coordinator is the leader of all the groups, while the *UAVs* are 'middle-managers'.

5.2 Minimal levels of group intention and awareness

What are the minimal levels of awareness and group intention needed for the agents on both inner- and outer-team levels?

5.2.1 The robots - two cases are applicable

1. They act only individually; this is the most limited (and economical) case;
2. They perform a limited form of cooperation, for example, they work together to clean up areas faster, or pitch in for other robots when these turn out to be unable to perform their part of the social plan.

The level of belief

1. In case 1, the robots need general belief about every group intention ($\text{E-BEL}_G(\text{E-INT}_G(\varphi))$) and about the distribution of the plan in bilateral commitments ($\text{E-BEL}_G(\bigwedge_{\alpha \in P} \bigvee_{i,j \in G} \text{COMM}(i, j, \alpha))$). This allows deliberation on actions of other robots and prevent them from doing all the work by themselves.
2. In case 2, E-BEL^2_G will be enough to allow deliberation about other robots' intentions and beliefs (especially $\text{E-BEL}^2_G(\text{E-INT}^2_G(\varphi))$. To see this, one may consider a pair of robots. With E-BEL^2_G, both robots have the same intention (this is the fact $\text{E-INT}_G(\varphi)$), believe they have the same intentions (the first-order belief $\text{E-BEL}_G(\text{E-INT}_G(\varphi))$), and believe that the other believes this (the second-order belief $\text{E-BEL}_G(\text{E-BEL}_G(\text{E-INT}_G(\varphi)))$). Therefore, the robots can reason about the beliefs and intentions of their partner.

The level of intention

1. In case 1, the robots need a general intention E-INT$_G$ about the goals.
2. In case 2, E-INT$_G^2$ will be enough to allow forming two-robot teams that are not competitive internally. (But see [6] for a counter-example showing that a two-level intention is not sufficient to preclude competition among two-agent coalitions). If agents are supposed to be strictly cooperative, a two-level definition is also sufficient for larger teams: all agents intend to achieve the goal with the others included in their team.

Although robots sometimes individually compete for resources, in our application where fast real-time team reaction to dangers is substantial, we opt for strictly cooperative robots that use fixed protocols to load up on resources. The cleanup robots do not communicate with robots from other teams, and therefore do not need to have any beliefs, intentions and commitments about them.

5.2.2 The *UAVs*

The *UAVs* must sometimes work with each other. This requires at least E-BEL$_G^2$ of other *UAVs*$'$ intentions.

The level of belief - Within each team of *UAV* and robots, the *UAV* has the highest level of awareness, and acts as a coordinator. In order to facilitate this (make plans and reason correctly), it will require one level of belief more than its agents:

- in case 1 we require $\text{BEL}(\mathit{UAV}, \text{E-BEL}_G(\text{E-INT}_G(\varphi)))$ with regard to the inner-team group intention $\text{E-INT}_G(\varphi)$ as well as: $\text{BEL}(\mathit{UAV}, \text{E-BEL}_G(\bigwedge_{\alpha \in \mathit{Cleanup}} \bigvee_{i,j \in G} \text{COMM}(i,j,\alpha)))$,
- in case 2 we require $\text{BEL}(\mathit{UAV}, \text{E-BEL}_G^2(\text{E-INT}_G^2(\varphi)))$ with respect to the level of inner-team group intention $\text{E-INT}_G^2(\varphi)$ as well as: $\text{BEL}(\mathit{UAV}, \text{E-BEL}_G^2(\bigwedge_{\alpha \in \mathit{Cleanup}} \bigvee_{i,j \in G} \text{COMM}(i,j,\alpha)))$.

The level of intention - Within the team, the *UAV* must make sure that all agents are motivated to do their tasks. Therefore:

- in case 1 we require $\text{INT}(\mathit{UAV}, \text{E-INT}_G(\varphi))$ with regard to the inner-team group intention $\text{E-INT}_G(\varphi)$,
- in case 2 we require $\text{INT}(\mathit{UAV}, \text{E-INT}_G^2(\varphi))$ with regard to the level of inner-team group intention $\text{E-INT}_G^2(\varphi)$.

5.2.3 The coordinator

The level of belief - One extra level of belief allows the coordinator introspection and reasoning about the joint effort of all *UAVs*. Therefore, since teams are cooperative in a limited way, we have $\text{BEL}(\mathit{coordinator}, \text{E-BEL}_G^{@@3@@}(\text{E-INT}_G^2(\varphi)))$ with respect to every group intention $\text{E-INT}_G^2(\varphi)$ as well as:
$\text{BEL}(\mathit{coordinator}, \text{E-BEL}_G^{@@3@@}(\bigwedge_{\alpha \in \mathit{Cleanup}} \bigvee_{i,j \in G} \text{COMM}(i,j,\alpha)))$.

The level of intention - Similarly, the coordinator has one level of intention more than the *UAVs* it manages, therefore we have $\text{INT}(\mathit{coordinator}, \text{INT}_G^{@@3@@}(\varphi)))$.

Commands from the coordinator overrule temporary contracts between teams. He does not only know the plan, but also keeps track of all relevant environmental conditions. We assume that even in the safe situation, the robots, *UAVs* and the pilot are prepared to take action at any moment.

5.3 Complexity of the language without collective attitudes

It seems that in the environmental case-study, the language used is richer than propositional modal logic. Fortunately, we can reduce most of the relevant part to a fixed finite number of propositional atoms (that may be combined and be the subject of attitudes), based on finitely many predicates and constants, as follows:

- a fixed number of relevant environmental states;
- a fixed number of pre-named locations;
- a fixed finite number of agents and teams;
- a fixed finite number of other objects (liquids, solids, catalyst, helicopter);
- a fixed number of relevant thresholds $n_1, n_2, n_3, \varepsilon_1, \varepsilon_2$.

The only possible source of unbounded complexity is the use of continuous intervals and real-valued functions $f_1, f_2, f_3, \mathit{fit}$. This can probably be simplified by using discretization.

6 Conclusion

In the case-study we have shown how to implement teamwork within a strictly cooperative, but still heterogenous group of agents in the TEAMLOG formalism. The heterogeneity is taken seriously here, as advocated in [13]. Natural differences in agents' shares, opportunities and capabilities when acting together, have been additionally reflected in different levels of agents' awareness about various aspects of their behaviour. The study dealt especially with cooperation and coordination. Having very generic definitions of common motivational and informational attitudes in TEAMLOG, it is challenging to choose a proper level of their complexity. We have shown that this is possible, by illustrating how to tailor complex definitions of intentions and commitments to a specific environment. For lack of space, not all the essential aspects of teamwork have been shown. Our focus was on building beliefs, intentions and, finally, commitments of all agents involved in teamwork on an adequate, but still minimal level. This way a bridge between theory and practice of teamwork has been effectively constructed for a specific application.

Future work will be to embed TEAMLOG into a form of approximate reasoning suitable for modeling perception, namely *similarity-based approximate reasoning*, which has intuitive semantics compatible with that of TEAMLOG [2, 5].

Acknowledgment

This research is supported by the Polish MNisW grant N N206 399334.

References

1. H. Aldewereld, W. van der Hoek, and J.-J. C. Meyer. Rational teams: Logical aspects of multi-agent systems. *Fundamenta Informaticae*, 63, 2004.
2. P. Doherty, B. Dunin-Kęplicz, and A. Szałas. Dynamics of approximate information fusion. In M. Kryszkiewicz, J. F. Peters, H. Rybinski, and A. Skowron, editors, *RSEISP*, volume 4585 of *Lecture Notes in Computer Science*, pages 668–677. Springer, 2007.

3. P. Doherty, G. Granlund, K. Kuchcinski, K. Nordberg, E. Sandewall, E. Skarman, and J. Wiklund. The WITAS unmanned aerial vehicle project. In *Proc. of the 14th European Conference on Artificial Intelligence*, pages 747–755, 2000.
4. P. Doherty, W. Łukaszewicz, A. Skowron, and A. Szałas. *Knowledge Representation Techniques. A Rough Set Approach*, volume 202 of *Studies in Fuzziness and Soft Computing*. Springer Verlag, 2006.
5. B. Dunin-Kęplicz and A. Szałas. Towards approximate BGI systems. In H.-D.Burkhard, G. Lindemann, R. Verbrugge, and L. Z.Varga, editors, *CEEMAS*, volume 4696 of *Lecture Notes in Computer Science*, pages 277–287. Springer, 2007.
6. B. Dunin-Kęplicz and R. Verbrugge. Collective intentions. *Fundamenta Informaticae*, 51(3):271–295, 2002.
7. B. Dunin-Kęplicz and R. Verbrugge. A tuning machine for cooperative problem solving. *Fundamenta Informaticae*, 63:283–307, 2004.
8. B. Dunin-Kęplicz and R. Verbrugge. Creating common beliefs in rescue situations. In B. Dunin-Keplicz, A. Jankowski, A. Skowron, and M. Szczuka, editors, *Proc. of Monitoring, Security and Rescue Techniques in Multiagent Systems (MSRAS)*, Advances in Soft Computing, pages 69–84, Berlin, 2005. Springer.
9. B. Dunin-Kęplicz and R. Verbrugge. Awareness as a vital ingredient of teamwork. In P. Stone and G. Weiss, editors, *Proc. of the Fifth Int. Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS'06)*, pages 1017–1024, New York (NY), 2006. IEEE/ACM Press.
10. M. Dziubiński. Complexity of the logic for multiagent systems with restricted modal context. In B. Dunin-Kęplicz and R. Verbrugge, editors, *Proc. of the Third Int. Workshop on Formal Approaches to Multi-Agent Systems, FAMAS'007*, pages 1–18. Durham University, 2007.
11. M. Dziubiński, R. Verbrugge, and B. Dunin-Kęplicz. Complexity issues in multiagent logics. *Fundamenta Informaticae*, 75(1-4):239–262, 2007.
12. R. Fagin, J. Halpern, Y. Moses, and M. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, MA, 1995.
13. N. Gold, editor. *Teamwork*. Palgrave McMillan, Basingstoke and New York, 2005.
14. J. Grant, S. Kraus, and D. Perlis. Formal approaches to teamwork. In S. Artemov and others, editors, *We Will Show Them: Essays in Honour of Dov Gabbay*, volume 1, pages 39–68. College Publications, London, 2005.
15. B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.
16. J. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37:549–587, 1990.
17. A. Kleiner, J. Prediger, and B. Nebel. RFID technology-based exploration and SLAM for search and rescue. In *Proc. of the IEEE/RSJ Int. Conference on Intelligent Robots and Systems (IROS 2006)*, pages 4054–4059, Bejing, 2006.
18. H. Levesque, P. Cohen, and J. Nunes. On acting together. In *Proc. Eighth National Conference on AI (AAAI90)*, pages 94–99, Cambridge (MA), 1990. MIT Press.
19. J.-J. C. Meyer and W. van der Hoek. *Epistemic Logic for AI and Theoretical Computer Science*. Cambridge University Press, Cambridge, 1995.
20. R. Parikh and P. Krasucki. Levels of knowledge in distributed computing. *Sadhana: Proc. of the Indian Academy of Sciences*, 17:167–191, 1992.
21. A. Rao and M. Georgeff. Modeling rational agents within a BDI-architecture. In R. Fikes and E. Sandewall, editors, *Proc. of the Second Conference on Knowledge Representation and Reasoning*, pages 473–484. Morgan Kaufman, 1991.
22. K. Sycara and M. Lewis. Integrating intelligent agents into human teams. In E. Salas and S. Fiore, editors, *Team Cognition: Understanding the Factors that Drive Process and Performance*, pages 203–232, Washington (DC), 2004. American Psychological Association.
23. M. Tambe. Teamwork in real-world, dynamic environments. In M. Tokoro, editor, *Proc. Second Int. Conference on Multi-Agent Systems*, pages 361–368, Menlo Park (CA), 1996. AAAI-Press.
24. M. Wooldridge and N. Jennings. Cooperative problem solving. *Journal of Logic and Computation*, 9:563–592, 1999.

Appendix

In all plans we assume we start from the base B where neutralizers are stored.

The social plan $\langle D_1 R\rangle$

Goal $\psi_1(L_1)$: to apply liquid L_1 on all areas contaminated with poison X_1.

{Assumption: One portion of L_1 neutralizes poison X_1 at a single location.}
while *contaminated-area* $\neq$ *emptyset* do
begin
  $A := calculate(UAV, \{rob_i\})$; {*UAV* finds region A for rob_i to clean up}
  $get(rob_i, L_1, B)$; {rob_i retrieves a tank with liquid L_1 from location B}
  $path := get\_path(UAV, rob_i, B, A)$; {$rob_i$ requests a path to follow}
  $move(rob_i, path)$; {rob_i moves from location B to location A}
  $pour(rob_i, L_1, A)$;
  contaminated-area := *contaminated-area* $\setminus A$;
  $return\_path := get\_path(UAV, rob_i, A, B)$;
  $move(rob_i, return\_path)$;
end.

The social plan $\langle D_1 N\rangle$

Goal $\psi_2(L_2)$: to spray liquid L_2 on areas contaminated with poison X_1.

{Assumption: One portion of L_2 neutralizes poison X_1 at a single location.}
{Assumption: The helicopter can transport k portions of liquid L_2.}
while *contaminated-area* $\neq$ *emptyset* do
begin
  $request(UAV, coordinator, pilot, \psi(L_2)$;
  $confirm(pilot, UAV, coordinator, \psi(L_2)$;
  $request(pilot, UAV, list_1, k)$;
  $send(UAV, pilot, list_1)$; {$list_1$ has at most k contaminated areas}
  $upload(helicopter, L_2)$; {*pilot* retrieves liquid L_2}
  take-off$(helicopter, B)$; {*pilot* takes off from location B}
  do $\langle$*plan-for-spraying*$(helicopter, L_2, l)\rangle$; {*pilot* sprays L_2 using his own invented plan}
  $confirm(pilot, UAV, done($*plan-for-spraying*$(helicopter, L_2, l))$;
  contaminated-area := *contaminated-area* $\setminus list_1$;
  $landing(helicopter, B)$;
  $free(pilot, coordinator)$;
end.

The social plan $\langle D_2 R\rangle$

Goal $\psi_3(S_1, K_1)$: to spread solid S_1 on all areas contaminated with poison X_2, followed by applying catalyst K_1 to all areas where S_1 is present.

{Assumption: One portion of S_1 and K_1 neutralize poison X_2 at a single location.}
while *contaminated-area* $\neq$ *emptyset* do
begin
  $A := calculate(UAV, \{rob_i, rob_j\})$;
  begin_parallel {operations are done in parallel}
    a plan similar to $\langle D_1 R\rangle$, but using S_1;
  ||
    $wait\_for(transporting(rob_i, S_1, A))$; {$rob_j$ waits until rob_i is on the way to A}
    $get(rob_j, K_1, B)$;
    $path := get\_path(UAV, rob_j, B, A)$;
    $move(rob_j, path)$;
    $wait\_for(spread(S_1, A))$; {rob_j waits for *someone* to spread S_1 in A}
    $pour(rob_j, K_1, A)$;
    $return\_path := get\_path(UAV, rob_j, A, B)$;
    $move(rob_j, return\_path)$;
  end_parallel
  contaminated-area := *contaminated-area* $\setminus A$;
end.

Plan $\langle D_2 N\rangle$ is similar to $\langle D_1 R\rangle$