DR 4.2: Models for roles and decisions in collaboration

Mario Gianni*, Manuel A. Ruiz Garcia*, Filippo Cantucci*, Federico Nardi*, Karthik Pushparaj† and Fiora Pirri*

*Alcor Laboratory, Department of Computer, Control, and Management Engineering "Antonio Ruberti" - Sapienza University of Rome.
†TU Dortmund, Germany.
⟨pirri@dis.uniroma1.it⟩

Project, project Id:	EU FP7 TRADR / ICT-60963
Project start date:	Nov 1 2013 (50 months)
Due date of deliverable:	M26
Actual submission date:	15th of February 2016
Lead partner:	ROMA
Revision:	final
Dissemination level:	PU

This document describes the progress status of the research on the development of models for roles and decisions in multi-collaboration. In this regard, we introduce a framework for managing the allocation of tasks of cooperative heterogeneous robots. Based on this framework, we model a multi-robot system as a set of multi-dimensional relational structures. These structures define collaborative tasks as both temporal and spatial relations between the processes and the tasks of the single robots. These relations are encoded in both a logical and geometrical fashion. We also describe a learning schema to extract from memorized experiences of collaborative task episodes the syntaxes, the semantics and the geometrical spaces in which the multi-dimensional relational structures lie. Moreover, we propose a decomposition technique for dealing with both the uncertainty of the data and the missing information as well as for knowledge discovering. Finally, we describe how both logical and the geometrical inference allows for task assignment. The document is organized as follows. Planned work is introduced and the

actual work is discussed, highlighting the relevant achievements, how these contribute to the current state of the art and to the aims of the project.

Executive Summary

The key objective of WP4 is to develop the formal methods needed to model knowledge exchange, knowledge maintenance, information sharing, common and individual decision structures in order to enable collaborative planning. In Year 2, we focused on the development of a framework for managing roles and decisions of robots, with particular attention to persistence. This framework develops persistence through learning. Both single- and multi-robot context-dependent behaviors are learnt from a collection of well-structured statements describing task episodes.

This framework has the main advantage to also learn those behaviors supporting collaboration which are not explicitly mentioned within mission reports. Moreover, it provides an inference mechanism for predicting the state of the robots. State prediction supports reasoning on roles and task assignment, in particular when communication between team members is interrupted due to WiFi signal loss. It also aims to reduce information exchange thus leading to a more careful use of the bandwidth of the communication channel (WP6).

Data collection, training and evaluation have been performed into a suitable Virtual Simulation Environment. This environment has been designed to simulate the dynamics of three robots, namely two UGVs and one UAV, similar to those of TRADR system. The main idea behind this work is to provide an environment for integrating functionalities of TRADR system. The virtual simulated environment also provides developers with a test-bed for conducting replicable testing of TRADR sub-systems in a decoupled manner (WP6).

Finally, a preliminary evaluation of performance of current approaches for task assignment in multi-agent planning has been conducted. This evaluation mainly focused on studying the features of planning frameworks accounting for private and public information exchange in order to reason about joint actions and plans.

Role of modeling decisions and role assignment in TRADR

Research work of WP4 in modeling decisions and roles assignment mainly contributes to Objective 2 of TRADR, as formulated in the Description of Work (DoW) of the project.

The key concept behind this objective is to develop persistence through the implementation of methods and models which allow a robot to learn, from its own experience and from the experience of other robots, gathered within and across sorties (and stored into the memory), how to better achieve goals in a previously unknown, hash environment.

The framework, described in Section 1.3.1, resorts to this key concept to develop persistence in modeling decisions and roles assignment, namely, learning from experience in memory. Indeed, it makes use of the information about robot-robot collaboration to learn multi-dimensional relational structures. The learned relational structures provide the basis of both a logical and a statistical inference mechanism which is responsible of dynamically regulating the assignment of a task to individuals or groups of robots. On the basis of both the learned relations and the data, describing task episodes, this mechanism provides estimates about what task a robot was performing, where the robot is supposed to operate, what is the failure rate of the task at hands, what is the event that induced the robot to switch to another task, what is the most probable reason why the robot does not communicate any-more with the remote command post, if a robot has to be supported by another robot in the neighbourhood in the execution of a task. These estimates support both human and robot decisions for the re-allocation of the tasks during mission execution (WP5).

Contribution to the TRADR scenarios and prototypes

Research work of WP4 in Year 2 provides the TRADR scenarios and prototypes with three main contributions: (1) the development of a framework for supporting decision-making for role and task assignment (WP5); (2) the development of a reasoning mechanism which does not require complete information of the state of the robots, thus reducing the amount of information which has to be exchanged through the communication channel (WP6) and, finally, (3) the development of an inference mechanism for retrieving the state of the robots when the communication channel is not reliable (WP6).

An additional contribution is provided by the designed Virtual Simulation Environment which offers a tool for both coupled and decoupled testing of TRADR functionalities for perception (WP1), planning and control (WP2). The possibility of building different simulated rescue scenarios makes this environment also suitable for use-cases specification as well as end-user training (WP3, WP5, WP7).

At this stage of the work, the proposed framework for role and task assignment is not fully integrated into TRADR system yet.

Persistence in WP4

Persistence in WP4 is addressed by developing a knowledge management structure for the team of TRADR robots whose key attribute is information reuse. Information reuse is meant from WP4 as (1) the access to the information gathered during previous operating activities and stored into the database; (2) the filtering of the information relevant for the tasks; (3) the integration of the filtered information with the incoming sensor measurements; (4) the reuse of the fused information for reducing the computational overhead in the current task assignment process in order to take into account the environmental changes and (5) learning from the registered experience in order to improve the performance of task execution in previously unknown environments.

In Year 2, WP4 mainly focused on the last aspect of aforementioned forms of information reuse, namely, on learning from experience. Indeed, we developed a framework capable of learning, from data reporting the situated history of the activities performed by the team of robots, a model of collaboration. From the data we extract information about the relations linking the robots, the tasks and the context features of collaborative tasks. This information is used to define both a set of symbols and terms of a many-sorted first-order language encoding multi-robot collaboration. This information is then used to learn the geometrical spaces underlying both individual and group task execution. These spaces are decomposed for obtaining the latent factors regulating the tasks of the robots. The decomposition is also used for dealing with the lack of information in the data as well as for discovering new forms of collaborations between robots.

1 Tasks, objectives, results

1.1 Planned work

The planned work of WP4, in Year 2, concerning "Models for roles and decisions in collaboration" is described in Task T4.2. Task T4.2 achieves the objectives described in Milestone MS4.2. An excerpt of the description of both Task T4.2 and Milestone MS4.2, from the DoW of the project, is given below

Task T4.2 *The goal of Task 4.2 is to extend the model defined in Task 4.1 so as to manage decisions and choices both in the generation of a common plan and during its execution, and maintain dynamic roles. In Task 4.2 knowledge flux is limited to the mission, to the robot team and only partially takes into account the knowledge that can be obtained from operators' requests. T4.2 expected result is to provide the models and the implementation to generate the common plan allocating roles and tasks, and execute it through the planned sorties. The plan is meant to work in a real environment, according to the effective integration with the other WPs, and in full within the augmented robot environment. Task 4.2 faces several challenges. To ensure continuity between generation and execution within collaboration, decisions in both plan generation and execution are modeled. This, in particular, requires to assign roles and tasks with time-space specification (though only within ranges) and with users commitments. Here collaboration is important in several phases and persistence at this stage is experimented with by preserving data and information across sorties along the whole time horizon of a mission.*

As specified in the early model for collaborative planning, described in Task 4.1, the basic knowledge and memory structure of each robot is designed to allot distinguished abilities to each of them, and to draw a net of constraints that is dynamically fitted to the collaborative plan. Constraints are used as meta-predicates to cope with the robot characteristics in the generation and distribution of tasks within the common goal. This basic structure, though elementary collaborative, thanks to the constraints on time, space and abilities, is still passive as it requires the operator to determine the common goal. To overcome this limitation, it is necessary that the robot team is able to generate sub-goals, distributing abilities, via roles allocation and sorties allocation. Our approach on role/task allocation extends the approaches based on both utility maximization (where the utility is defined in terms of energy consumption, computational requirements, resources, probability of success) and auction with a learning schema on performance of individuals and groups across sorties.

To predict both group performance, and time performance for specific tasks, depending on context features, such as amount of rubble, light condi-

tions, smog, dust, stairs, communication availability, and similar, relations between the robots, the tasks and the context features have to be specified. The entities got involved in the relations are represented by categorical variables. On the other hand, the relations between categorical variables are specified in terms of tensor fields. Techniques based on Tensor Decomposition are applied for obtaining the latent factors describing the relations among the categorical variables [1, 36, 12]. Due to the presence of few experiments not covering all possible group compositions and assuming that the latent factors can change across different sorties, a Bayesian approach to tensor factorization is taken to model the evolution of latent factors [59]. The outcome is a recommendation list for group formation, for time horizon for sorties, and for information sharing. As specified in Task 4.1 the model requires a precise coding of terms and features into structured memories. Inference on roles and tasks might need to be revised during sorties, because of stimuli and events that change the context; here learned recommendations are used, also for updating in shifting tasks.

Milestone MS4.2 *MS4.2 proves the models defined in Task 4.2 extending the early collaborative planning designed in Task 4.1. More specifically, given a mission composed of a number of sorties, it proves task allocation models at different stages of the mission. It proves decision models for both plan generation and during task execution in each sortie, which might need a reallocation due to failures. It tests time horizon assignment, within ranges, for each sortie and the scheduling of information that has to be sent to the operators. It shows when and how decisions for the area each robot want/can cover, are taken by reasoning on the representation of the integrated map, provided by the perception functionalities. Finally, it proves robot reasoning methods about own task processes, in relation to the others, and robot reasoning methods about others task processes with respect to its own actions.*

1.2 Addressing reviewers' comments

We have addressed the reviewers' comments as follows:

1. *The teamwork resiliency issue should be carefully addressed so that the multi-robot system can take advantage from its inherent redundancy in the presence of failures of individual robots.*

   The proposed framework for modeling decisions and role assignment is endowed with a mechanism for predicting robot as well as task execution state (see Actual work performed section, Sub-section 1.3.1). This mechanism turns out to be quite useful, especially when robot status can not be monitored due to failures, such as wifi signal lost. According to the predicted state, the framework can choose among several operational control strategies and recovery procedures to support both individual and collective task execution. This approach increases both robustness and reliability of the multi-robot system.

2. *The motivation for proposing a new multi-robot collaboration method should be better justified by comparing it with state-of the art methods on multi-robot cooperation.*

   The motivations for proposing an alternative approach are numerous and concern the core of TRADR project, that is, the development of persistence.

   Persistence is strictly related to the concept of memory and to the re-use of information stored therein. Memory has to be structured to filter information from noise, to facilitate the access to the information and to reduce information overload. On the other hand, information re-use includes not only querying the data loaded in memory but also learning. The main objective of learning is twofold: to learn behavior rules of both individual and cooperative robots and, secondly, to fill the empty spaces of the memory due to missing information, namely, to discover new information. By filling the memory, it can be possible to retrieve estimates of the state of events, through both logical and statistical inference, without the need of explicit knowledge about the current state of the robots as a whole. This implies that either a robot or an operator, depending on the physical allocation of the memory (e.g., centralized or distributed), does not need to make a request to other members in order to take a decision. A team member can make use of the estimates of the events to decide what to do next.

   To date, as far as we can tell, there are no approaches that fully develop this paradigm in an unified framework. Therefore, it turns out to be quite difficult to make a comparison. However, an evaluation has been performed. In this evaluation we considered current state-of-the-art planning frameworks for cooperative multi-robot systems which account for the problem of the choice of the information that has to be sent through the communication channel in order to reason about joint actions [5, 7, 19, 65, 66]. The algorithms that we considered aim to deal with this issue whilst solving planning problems.

   In Actual work performed section, Sub-section 1.3.1, we clarify the main concepts and ideas behind the framework for modeling decisions and role assignment. In Annexes, Section 2.2, the evaluation results are reported in the context of TRADR scenarios and prototypes

3. *An experimental evaluation study based on a quantitative comparison*

of the proposed method with state-of-the-art methods is recommended.

Actually, making a quantitative comparison with state-of-the-art methods is quite difficult. In fact, most task-related performance measures, like the execution time, rather describe the performance of external factors like the software/hardware needed for executing robots than the framework developed for representing the multi-robot system. Comparison is also difficult due to the lack of approaches which implement persistence linking together memory, learning and knowledge discovery. However, general qualitative criteria can be formulated in order to perform an evaluation of the framework together with a comparison with state-of-the-art methods. These criteria account for the expressiveness of the framework encoding the multi-robot system, the capability of the framework of supporting reasoning about task assignment for a wide class of collaborative tasks, the capability of the framework of supporting reasoning in the presence of either communication failures or of missing information. A more detailed comparison, on the basis of these criteria, will be definitely our main objective in the nearby future.

4. *Surprisingly, two different but related path- planning methods were presented in WP2 (D*-Lite) and in WP4 (A*), the partners involved in these WPs should carry out a comparative study of those algorithms in order to choose one single algorithm to be used afterwards in the project.*

   We modified the Description of Work (DoW) of TRADR project in order to take the main responsibility of all the aspects concerning persistence in both individual and collaborative path planning of the UGVs. Research work of WP4 regarding planning is reported in Actual Work Performed section, Task 2.5, in Deliverable DR2.2.

5. *A joint paper of the involved partners to present this study is recommended.*

   We were not able to follow up on this recommendation. We decided to focus our efforts on the achievement of Task 2.5, reported in Deliverable DR2.2 as well as toward the development of the proposed framework for roles and task assignment, for which an article for the Artificial Intelligence Journal is in preparation.

6. *In WP1 and WP4, special attention should be given in the future to communication contention methods when scaling to larger teams of robots.*

   Actually this is a crucial problem in multi-robot systems. Task assignment requires knowledge to be exchanged among the members of a team. However, this knowledge exchange overloads the communication channel as the number of units involved in a rescue mission increases. Although this issue will be faced by WP4 in Year 3, in Year 2, an approach to reduce the communication overhead is proposed. This approach relies on the prediction mechanism of the proposed framework to retrieve the information needed for assigning roles and tasks to the robots, without resorting to explicit message passing mechanisms.

7. *It is not clear how human interaction will be integrated in multi-robot collaboration. This is another open issue that should be addressed in the next year together with WP5.*

   Interaction with humans is going to be faced by both WP4 and WP5 in Year 3. However, in Year 2 basic operational modes of interaction are already provided. First mode concerns the possibility of an operator to interact with a team of robots by posting goals and objectives. Second mode is to provide operators with feedback about the decisions on the tasks assigned tothe robots, given the goals, as well as with the possibility to change these roles, before and during execution.

1.3 Actual work performed

The actual work performed supports the objectives of Milestone MS4.2. This work focused on the development of a framework for modeling the decision-making processes responsible for task assignment in collaborative heterogeneous multi-robot systems. The framework first builds a set of multi-dimensional relational structures modeling the collaboration on the basis of both prior knowledge and a learning schema. Then, it develops the mechanisms underlying the decision-making processes for task assignment. On the basis of the learned structures, these mechanisms rely on two different models of inference, which combined together allow for reasoning about collaborative tasks. The first model uses logical entailment. At this stage, task assignment is addressed as a deduction problem. However, this inference appears to be very limited in multi-robot domain applications, due to both incompleteness of the prior knowledge and sparsity of information after learning. So this is where the second model of inference comes into play. This model exploits the quantitative representation of the collaboration among heterogeneous robots, provided by the multi-dimensional relational structures, to reason about multiple alternatives of cooperation in the presence of both incomplete and missing knowledge. Knowledge that logical inference lacks is discovered by this model through prediction. New links between individual tasks of single robots, not modeled in the prior knowledge as well as not learnt due to missing data, are predicted through a geometrical decomposition of the collaborative tasks. Learning of the structures supporting decision-making together with knowledge discovery through link prediction are crucial steps toward the development of persistence in multi-robot collaboration.

In Year 2 WP4 also devoted time to the development of a Virtual Simulation Environment, where data collection for both building and learning of the relational structures underlying collaboration, learning itself and testing took place.

In the following, this section reports, the research carried out by WP4, as briefly introduced above, contributing to Task T4.2.

1.3.1 Role and task allocation framework for Multi-Robot Collaboration with latent knowledge estimation

In the proposed framework, we modeled a team of heterogeneous cooperative robots as a set $\mathfrak{U}$=$\{\mathcal{S}_1,\ldots,\mathcal{S}_n\}$ of multi-dimensional relational structures. Each structure $\mathcal{S}_i$ comprises a many-sorted signature $\mathfrak{S}_i$ and a non-negative multi-dimensional matrix $\underline{\mathbf{Y}}_i$, also called tensor. Each signature $\mathfrak{S}_i$ includes

- a relational symbol R_i of arity $K\in\mathbb{N}$;
- a finite set $\Sigma_i=\{\sigma_1,\ldots,\sigma_{n_i}\}$ of types, also called sorts;
- a finite set C_{i,σ_k} of constant symbols c_{i,σ_k}, for each sort $\sigma_k\in\Sigma_i$;
- a countable set V_{i,σ_k} of variable symbols v_{i,σ_k}, for each sort $\sigma_k\in\Sigma_i$.

Each tensor $\underline{\mathbf{Y}}_i\in\mathbb{R}_+^{M_1\times\cdots\times M_K}$ has

- a number of dimensions equal to the arity K of R_i;
- a number of elements, along the k-th direction, equal to the cardinality M_k of the set C_{i,σ_k}, where $\sigma_k\in\Sigma_i$ is the sort of the k-th input argument of R_i.

According to the definition of $\underline{\mathbf{Y}}_i$, each tuple of indices $(i_1,\ldots,i_K)$, with $i_k=1,\ldots,M_k$, corresponds to a tuple $(c_{i,\sigma_1},\ldots,c_{i,\sigma_K})$ of constant symbols, with $c_{i,\sigma_k}\in C_{i,\sigma_k}$ and $\sigma_k\in\Sigma_i$ sort of the k-th input term of R_i, for $k=1,\ldots,K$.

In the following we illustrate with an example how the multi-dimensional relational structures, as defined above, encode forms of cooperation among robots.

Example 1.1. Let us consider a multi-robot system composed of two heterogeneous robots, that is, an UGV, named `UGV1`, and an UAV, named `UAV2`. Suppose that `UGV1` has to explore an area of the environment, e.g., the first floor of a collapsed building (see Figure 1(a)). In order to navigate, this robot needs to have a representation of the area specifying what regions are traversable. The analysis of traversability builds upon a 3D metric representation of the area. However, it might happen that, due to the high degree

(a) UGV1 exploring first floor of a simulated collapsed building.

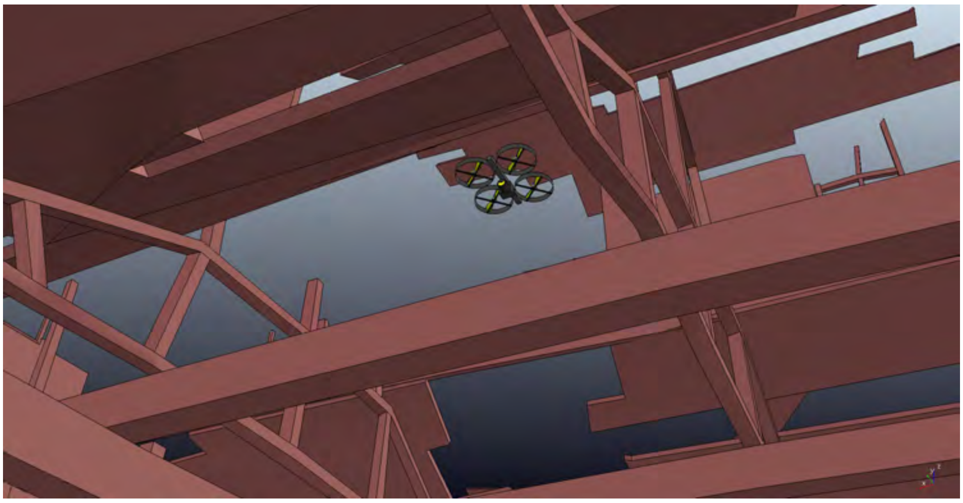

(b) UAV2 flying over the roof of the simulated collapsed building.

Figure 1: UGV1 and UAV2 in the virtual simulated environment.

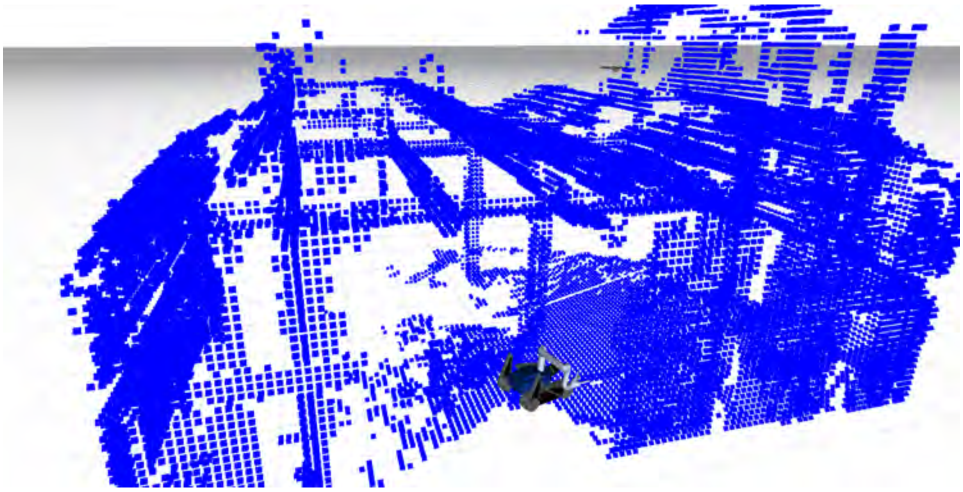

(a) 3D metric map of the area built by UGV1.

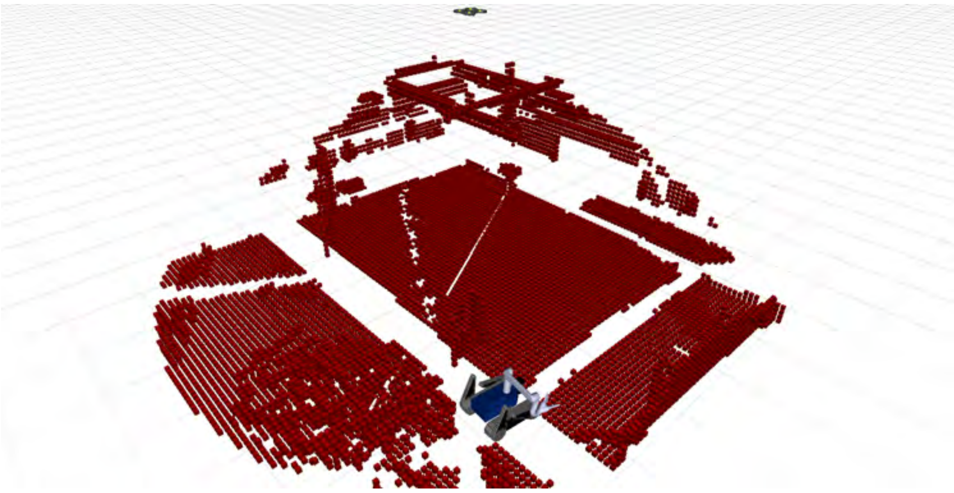

(b) 3D metric map of the area built by UAV2.

Figure 2: ROS RVIZ visualization of both the UGV1 and UAV2 inside the simulated collapsed building.

of harshness of the terrain, the metric map built by UGV1 is quite sparse, thus making traversability analysis very inaccurate (see missing points on the left side of the robot in Figure 1(c)). Under this situation, UGV1 sends the request to UAV2 to fly over the area to build a more dense metric map (see Figure 1(b) and (d)). Upon the completion with success of this process, UGV1 requests the map of the area to UAV2, it integrates its own map with the map provided by UAV2 and, finally, it computes a more accurate estimate of the traversability of the surrounding. This form of collaboration can be represented by a multi-dimensional relational structure $\mathcal{S}$ encoding a temporal relation Before between four entities, two of type Robot and two of type Process, whose signature $\mathfrak{S}$ is defined as follows

$$K = 4, \quad \Sigma = \{\texttt{Robot}, \texttt{Process}\}, \quad R(\cdot,\cdot,\cdot,\cdot) \stackrel{\text{def}}{=} \texttt{Before}(\cdot,\cdot,\cdot,\cdot),$$
$$C_{\texttt{Robot}} = \{\texttt{UAV2}, \texttt{UGV1},\ldots\} \quad \text{and} \quad C_{\texttt{Process}} = \{\texttt{mapping}, \texttt{exploring},\ldots\}.$$

Finally, for the tuple ($\texttt{UAV2}, \texttt{mapping}, \texttt{UGV1}, \texttt{exploring}$) of constant symbols, there exists a tuple of indices (i_1, i_2, i_3, i_4) such that $y_{i_1,i_2,i_3,i_4} \in \mathbb{R}_+$ is the element of the four-tensor $\underline{\mathbf{Y}}$ associated with this tuple.

In the next example we describe how multi-dimensional relational structures are used for reasoning about role and task assignment.

Example 1.2. Let us consider a multi-robot system composed of three UGVs, named UGV1, UGV2 and UGV3, respectively. Let us suppose that $\mathfrak{U}$ includes a structure $\mathcal{S}_1$, having the following signature $\mathfrak{S}_1$

$$K = 4, \quad \Sigma_1 = \{\texttt{Robot}, \texttt{Process}\}, \quad R_1(\cdot,\cdot,\cdot,\cdot) \stackrel{\text{def}}{=} \texttt{Equal}(\cdot,\cdot,\cdot,\cdot),$$
$$C_{1,\texttt{Robot}} = \{\texttt{UGV1},\texttt{UGV2},\texttt{UGV3},\ldots\} \quad \text{and} \quad C_{1,\texttt{Process}} = \{\texttt{grasping},\ldots\}.$$

$\mathcal{S}_1$ represents a collaborative pick and place task. This task requires that two UGVs simultaneously hold and lift an object. According to this specification, this form of collaboration has been encoded by a temporal relation Equal between four entities, two of type Robot and two of type Process. Moreover, let us assume to have another structure $\mathcal{S}_2$ representing the processes which can be performed by each robot. The signature $\mathfrak{S}_2$ of $\mathcal{S}_2$ is defined as follows

$$K = 2, \quad \Sigma_2 = \{\texttt{Robot}, \texttt{Process}\}, \quad R_2(\cdot,\cdot) \stackrel{\text{def}}{=} \texttt{Process}(\cdot,\cdot),$$
$$C_{2,\texttt{Robot}} = \{\texttt{UGV1},\texttt{UGV2},\texttt{UGV3},\ldots\} \quad \text{and} \quad C_{2,\texttt{Process}} = \{\texttt{grasping},\ldots\}.$$

Now, suppose that UGV1 has to grasp an object for which it is required the collaboration of another UGV. Then, we want to know to which robot to assign this collaborative task.

Modeling a multi-robot system via multi-dimensional relational structures allows for two different type of inference. The first type is the standard logical entailment. The second is based on the extraction of fragments of the tensors through modes-n operations. By combining both logical inference and reasoning on tensors a decision-making process for task assignment is devised, as described below.

According to logical entailment, we know that, both `UGV2` and `UGV3` can execute this task. Now, let $y^1_{i_1,i_2,\bar{i}_3,i_4} \in \mathbb{R}_+$ and $y^1_{i_1,i_2,\hat{i}_3,i_4} \in \mathbb{R}_+$, at positions $(i_1, i_2, \bar{i}_3, i_4)$ and $(i_1, i_2, \hat{i}_3, i_4)$, be the elements in the four-tensor $\underline{\mathbf{Y}}_1$ associated to the following tuples

$$
\begin{aligned}
y^1_{i_1,i_2,\bar{i}_3,i_4} &\longmapsto \langle \texttt{UGV1}, \texttt{grasping}, \texttt{UGV2}, \texttt{grasping} \rangle \\
y^1_{i_1,i_2,\hat{i}_3,i_4} &\longmapsto \langle \texttt{UGV1}, \texttt{grasping}, \texttt{UGV3}, \texttt{grasping} \rangle
\end{aligned}
$$

Let us also assume that these elements represent the number of times that `UGV1` performed this task in collaboration with the two UGVs under consideration. According to the meaning that we are given to these tensor elements, the choice of the UGV which has to collaborate with `UGV1` might be dictated by comparing the value of $y^1_{i_1,i_2,\bar{i}_3,i_4}$ with the value of $y^1_{i_1,i_2,\hat{i}_3,i_4}$. Suppose that $y^1_{i_1,i_2,\hat{i}_3,i_4} > y^1_{i_1,i_2,\bar{i}_3,i_4}$. Then the choice falls onto `UGV3`.

As above, let $(\bar{i}_1, i_2)$, $(\hat{i}_1, i_2)$ and $(\tilde{i}_1, i_2)$ be the three tuples of indices such that $y^2_{\bar{i}_1,i_2} \in \mathbb{R}_+$, $y^2_{\hat{i}_1,i_2} \in \mathbb{R}_+$ and $y^2_{\tilde{i}_1,i_2} \in \mathbb{R}_+$ are the elements in the matrix $\mathbf{Y}_2$ associated with the terms in the interpretation of `Process`

$$
\begin{aligned}
y^2_{\bar{i}_1,i_2} &\longmapsto \langle \texttt{UGV2}, \texttt{grasping} \rangle \\
y^2_{\hat{i}_1,i_2} &\longmapsto \langle \texttt{UGV3}, \texttt{grasping} \rangle
\end{aligned}
\tag{1}
$$

Now, suppose that these elements $y^2_{\hat{i}_1,i_2}$ and $y^2_{\tilde{i}_1,i_2}$ represent the failure rate of the process `grasping` for `UGV2` and `UGV3`, respectively. A comparison among the values of these two tensor elements should be also considered before assigning the task to `UGV3`. If the value of $y^2_{\tilde{k}_1,k_2}$ is greater than the value of $y^2_{\hat{k}_1,k_2}$ then a better choice would be to assign the collaborative task to `UGV2`, rather than `UGV3`.

We developed a schema through which both signatures and tensors of each multi-dimensional relational structure are learnt from data reporting the situated history of the activities performed a group of cooperative heterogeneous robots. Here, the schema exploits the linguistic structure of the data to make learning more tractable.

In this regard, data describing reports of missions executed by a team of robots are organized as a treebank [22, 31, 40, 63]. A treebank is a collection of pairs $\langle \mathbf{s}_i, \mathcal{T}_i \rangle$, where each $\mathbf{s}_i$ is a statement and each $\mathcal{T}_i$ is a syntactic tree. Each statement is a sequence $s_{i,1}, \ldots, s_{i,n}$ of words. Here, we assume that

Σ_i={Robot, Stimulus, Process}
$R_i(\cdot,\cdot,\cdot) \overset{\text{def}}{=}$ Inhibition$(\cdot,\cdot,\cdot)$
$C_{i,\texttt{Robot}}$={r5,...}
$C_{i,\texttt{Stimulus}}$={wifi connection lost,...}
$C_{i,\texttt{Process}}$={moving,...}

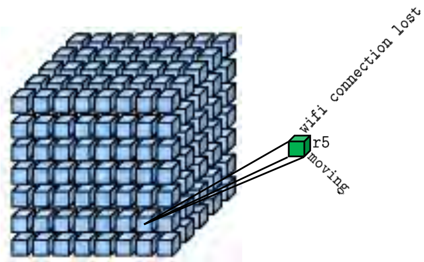

Figure 3: Signature and the 3-order tensor associated with the three-dimensional relational structure representing inhibition behaviors of a robot, with respect to both processes and stimulus occurrences.

statements do not contain anaphoric as well as elliptical references. Each $\mathcal{T}_i$ is composed of a set $\mathcal{N}_i$ of nodes and a set $\mathcal{E}_i$ of edges. $\mathcal{N}_i$ is composed of a set $\mathcal{N}_{i,int}$ of intermediate nodes and a set $\mathcal{N}_{i,leaf}$ of leaf nodes. Each node $n_u \in \mathcal{N}_{i,int}$ is labeled with a non-terminal symbol NT of a formal system FS. On the other hand, each node $n_v \in \mathcal{N}_{i,leaf}$ is labeled with a terminal symbol T of FS. There exists an edge $\langle n_u, n_v \rangle \in \mathcal{E}_i$, with $n_u, n_v \in \mathcal{N}_i$ if there exists a production rule of the form $\gamma \to \delta$ in FS such that the label of n_u is equal to γ, the label of n_v is equal to δ, $\gamma \in NT$ and $\delta \in NT \cup T$. Intuitively, each syntactic tree is a derivation of the string of words composing a statement. A derivation is a sequence of rule expansions defined by the formal system.

This system provides *a priori* a well-defined syntax of the relations encoding both individual and collaborative tasks. It has the effect to filter noise in the documents. Moreover, it provides the semantics of the relations thus speeding up learning of the structures needed for modeling the multi-robot system. More details about the specification of the formal system are described in Annexes 2.1.

The learning schema comprises three main steps: (1) the definition of the signatures; (2) building of the tensors and, finally, (3) the estimation of the values of the elements of each tensor.

Signatures are directly derived from the specification of both constituents and production rules of the formal system and from the syntactic trees annotating the statements. Given the signatures, extracted from the treebank, we build the tensors of the structures on the basis of their definition. More precisely, the number of dimensions of $\underline{\mathbf{Y}}$ is fixed to be equal to the arity K of $R \in \mathcal{S}$. If the k-th input term of R is of sort $\sigma_k \in \Sigma$ then, along k-th direction the number of elements is fixed to be equal to the cardinality M_k of the constant set C_{σ_k}. Moreover, each index $i_k \in \{1,\ldots,M_k\}$ is linked to one and only one constant symbol $c_{\sigma_k} \in C_{\sigma_k}$. The 3-order tensor associated with the three-dimensional relational structure representing inhibition behaviors of a robot, on the basis of these rules, is illustrated in Figure 3. Finally,

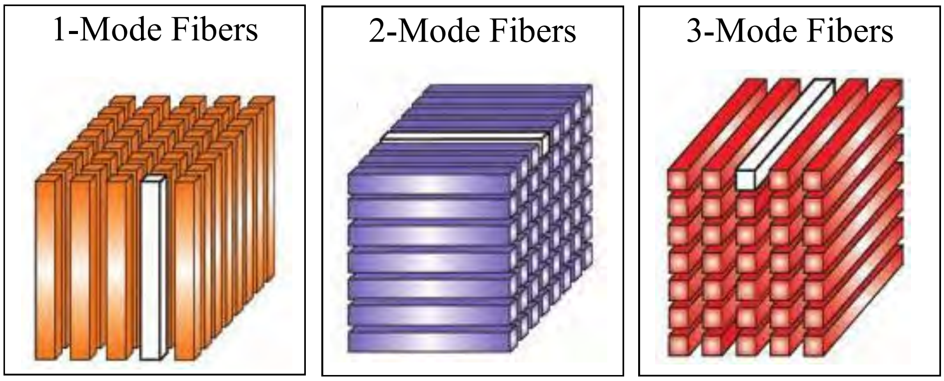

Figure 4: From left to right, 1-mode, 2-mode and 3-mode fibers of a third-order tensor.

the last step of learning of the multi-dimensional relational structures is to estimate the values of the entries of the tensors. To this end, we resorted to a generative approach based on a probabilistic model, similar to those applied for document classification and information retrieval [51, 54, 62]. This approach is detailed in Annexes 2.1.

However, after this step, the tensor turns out to be sparse due to the lack of statements in the treebank. In order to fill the missing entries of the tensors we resorted to Non-Negative Tensor Decomposition (NTD) [1, 36, 12]. NTD reduces the dimensionality of the tensor through factorization. The components resulting from this factorization require less computational resources for both storage and information retrieval. NTD also filters the data thus reducing noise. Finally, by applying NTD, new knowledge can be discovered through link prediction [23, 60].

Approaches based on Alternating Least Squares (ALS) minimization of the squared Euclidean distance are commonly employed for estimating the components in NTD [1]. However, these approaches do not deal with missing entries in $\underline{\mathbf{Y}}$. In order to cope with this issue, we proposed an approach which embeds a variant of ALS, named Fast Hierarchical Alternating Least Squares (F-HALS) [49], into an imputation-alternation schema [43]. Details about this approach are provided in Annexes 2.1.

Once the components have been estimated through this algorithm mode-k fiber operations can be performed for extracting one-dimensional fragments of the tensors. Fibers are obtained by fixing all indexes except one, as illustrated in Figure 4. mode-k fiber operations can be applied for reasoning about multiple choices in a multi-robot collaboration setting. In Example 1.3 we illustrate how these operations support reasoning about task allocation as well as how link prediction leads to knowledge discovery.

Example 1.3. Let us consider the scenario described in Example 1.2. `UGV1` has to grasp an object for which it is required the collaboration of another UGV. Therefore, we have to take a decision regarding which UGV has to take at hands this collaborative task. From logical inference we know that both `UGV2` and `UGV3` can execute this task. Now, let us assume that there exists another another UGV, namely `UGV4`, which might handle this task. Let us also assume that there was no statement in the treebank reporting a collaboration of this form between `UGV1` and `UGV4`. Under this perspective, the learning schema has not been able to derive the instance $\langle \texttt{UGV1}, \texttt{grasping}, \texttt{UGV4}, \texttt{grasping} \rangle$ of the relation `Equal` between entities of sort `Robot` and of sort `Process`.

However, according to the procedure with which we have built the tensor associated with this multi-dimensional structure, there exists a third tuple of indices $(i_1, i_2, \tilde{i}_3, i_4)$ such that $\hat{y}^1_{i_1,i_2,\tilde{i}_3,i_4}$ is the estimated value of tensor element, resulting from the decomposition of $\underline{\mathbf{Y}}_1$, assigned to the term $\langle \texttt{UGV1}, \texttt{grasping}, \texttt{UGV4}, \texttt{grasping} \rangle$, even if this term is not an instance of the relation `Equal`. In other words, through the decomposition, we can discover, with a certain degree of confidence, new knowledge. If we suitably perform a fiber operation on the tensor then we now obtain a one-dimensional fragment whose entries are an estimation of the collaboration of `UGV1` with all the other UGVs (or at least with all included in the domain of discourse). By interpreting this fragment as a recommendation vector, we can take a decision about which UGV has to be in charge of supporting `UGV1` in the collaborative task [2].

From a different point of view, mode-k operations on tensors provide estimates of instances of both individual and collaborative behaviors of the robots. We exploited this important feature of the framework for dealing with situations in which re-allocation of tasks is required but failures of the communication channel hinder the information exchange needed for taking a decision. In this regard, the proposed framework allows us to estimate, in the case in which a robot does not communicate with the remote command post, what task the robot was performing, what was the main cause of the interruption of the communication, what the robot has decided to do on the basis of the occurrence of an event such as communication failure due to WiFi signal lost and if this robot needs to be supported by another robot or not for the accomplishment of the task at hands.

Several experiments have been carried out in the designed Virtual Simulated Environment to qualitatively evaluate the main features of the framework we devised for role and task assignment in multi-robot collaboration. Details about the obtained results are discussed in Annexes 2.1.

1.4 Relation to the state-of-the-art

In this section we describe how the results of Deliverable DR4.2 relate to the state-of-the-art.

Multiple-robot systems (MRS) for Urban Search&Rescue (USAR) domain applications develop advanced robotic technologies to improve the effectiveness of multiple robot deployment in disaster response [61]. Deployment of multiple robots in rescue scenarios increases area coverage thus increasing the chance to find all potential survivors [9]. Homogeneous robot deployment introduces redundancy making the overall system more robust and reliable [47]. Conversely, the use of heterogeneous robots can cope with hardware limitations of individual robot payload for mission accomplishment [27, 3, 28]. Moreover, multiple robots deployment minimizes human exposure to danger [10, 11].

MRS have received significantly increasing attention over the last decades [48]. Several research efforts in robotics have been made in order to effectively develop MRS for safety-critical domain applications, mainly focusing on organization [41, 18], collaboration [56, 70], coordination [33, 55], task allocation [13, 20, 30], negotiation [58], communication [67, 4], team composition [29], team performance [68, 32], architectures [46, 45, 39], coalition [69, 38, 72], control [8, 42, 73], fault detection [15, 35] and human-robot interaction [14].

Our research work on modeling roles and task assignment is mainly related to the research field in MRS which concerns with knowledge management and information sharing [71, 26, 16, 21, 64]. Frameworks for developing MRS endowed with a structure for knowledge management typically provide a language for representing knowledge, a mechanism for knowledge association, a language for communicating knowledge and, finally, a memory system. Languages for representing knowledge in MRS are commonly based on beliefs intentions [53], semantic networks [37], frame languages [71] and resource description frameworks [6, 34, 52]. Knowledge association is responsible of the bidirectional information flow, where low-level data is passed upwards and the high-level information is returned downwards using logical inference [34], bayesian inference [53], semantic relationships and hierarchies [57] or computational learning methods [50]. FIPA [44] together with KIF [25] are the standard languages for communication. Finally, memory is usually deployed on either centralized [17] or on distributed [24] database systems.

To date, very few MRSs integrate all the aforementioned components in a common framework supporting abstract reasoning [64]. Moreover, we are not aware of any other MRS that learns from data stored within the memory system. Conversely, we proposed a preliminary approach that exploits data collected about environments, objects and action log data to enable the MRS to learn typical event occurrences, success models of tasks given the context, common execution failures, timing information, or promising collaborative action for a given group of heterogeneous robots. Besides learning, we also focused on proposing a preliminary solution, based on prediction, to the problem of information exchange among robots in the presence of failures of the underlying communication infrastructure [4]. Moreover, the proposed framework for modeling roles and task assignment in multi-robot collaboration extends current state-of-the-art on knowledge representation and reasoning by devising a method for knowledge discovery.

2 Annexes

2.1 Gianni, Pirri(2015), "Roles and task allocation framework for Multi-Robot Collaboration with latent knowledge estimation"

Bibliography M. Gianni, F. Pirri. "Roles and task allocation framework for Multi-Robot Collaboration with latent knowledge estimation". Technical report. Alcor Laboratory, DIAG "A. Ruberti", Sapienza University of Rome.

Abstract In this work a novel framework for roles and task allocation in Cooperative Heterogeneous Multi-Robot Systems (CHMRSs) is presented. This framework models a CHMRS as a set of multi-dimensional relational structures (MDRSs). This set of structure defines collaborative tasks as both temporal and spatial relations between processes and tasks to be performed by a team of heterogeneous robots. These structures are enriched with tensor fields which allow for geometrical reasoning about collaborative tasks. A learning schema is also proposed in order to derive the components of each MDRS. This schema relies on three main key aspects: (1) the definition of a precise linguistic structure, called Multi-Robot Collaboration Treebank (MRCT) constraining the data used for building the multi-dimensional relations, (2) the application of a generative approach based on a probabilistic model for building the tensor fields and, finally, (3) the estimation of latent knowledge through Non-Negative Tensor Decomposition (NTD). Preliminary evaluation of the performance of this framework is also presented in simulation with three heterogeneous robots, namely, two Unmanned Ground Vehicles (UGVs) and one Unmanned Aerial Vehicle (UAV).

Relation to WP This work contributed to the development of a framework for multi-robot collaboration focusing on persistence for the accomplishment of T4.2.

Availablity Unrestricted. Included in the public version of this deliverable.

2.2 Gianni, Ruiz, Pirri(2015), "Evaluation of Multi-Agent Planning frameworks under distributed privacy-preserving"

Bibliography M. Gianni, M. A. Ruiz Garcia, F. Pirri. "Evaluation of Multi-Agent Planning frameworks under distributed privacy-preserving". Technical report. Alcor Laboratory, DIAG "A. Ruberti", Sapienza University of Rome.

Abstract In this report we describe the evaluation of current state-of-the-art Multi-Agent planning frameworks focusing on privacy-preserving. In these frameworks, agents can decide to share public capabilities while keeping private processes and information that support these capabilities. Several instances of well-known cooperative multi-agent problems have been considered for this evaluations. Plan quality, planning time, ratio of problem solved and dimension of public information have been used for measuring performance of the planners under consideration. Results and discussion are also reported, in particular, under the point of view of TRADR scenarios and prototypes.

Relation to WP This work contributes to Task T4.2 in evaluating performance of state-of-the-art Multi-Agent planning frameworks accounting for distributed and parallel planning, planning reuse and information sharing.

Availablity Unrestricted. Included in the public version of this deliverable.

References

[1] E. Acar and B. Yener. Unsupervised multiway data analysis: A literature survey. *IEEE Transactions on Knowledge and Data Engineering*, 21(1):6–20, 2009.

[2] Gediminas Adomavicius and Alexander Tuzhilin. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Trans. on Knowl. and Data Eng.*, 17(6):734–749, 2005.

[3] S. Balakirsky, S. Carpin, A. Kleiner, M. Lewis, A. Visser, J. Wang, and V. A. Ziparo. Towards heterogeneous robot teams for disaster mitigation: Results and performance metrics from robocup rescue: Field reports. *Journal of Field Robotics*, 24(11-12):943–967, 2007.

[4] S. Blumenthal, N. Hochgeschwender, E. Prassler, H. Voos, and H. Bruyninckx. An approach for a distributed world model with qos-based perception algorithm adaptation. In *IROS*, pages 1806–1811, 2015.

[5] Daniel Borrajo. Multi-agent planning by plan reuse. In *AAMAS*, pages 1141–1142, 2013.

[6] A. Bouguerra, L. Karlsson, and A. Saffiotti. Semantic knowledge-based execution monitoring for mobile robots. In *ICRA*, pages 3693–3698, 2007.

[7] Ronen Brafman. A privacy preserving algorithm for multi-agent planning and search. In *IJCAI*, 2015.

[8] B. Browning, J. Bruce, M. Bowling, and M. M. Veloso. Stp: Skills, tactics, and plays for multi-robot control in adversarial environments. *Proceedings of the Institution of Mechanical Engineers, Part I: Journal of Systems and Control Engineering*, 219(1):33–52, 2005.

[9] W. Burgard, M. Moors, C. Stachniss, and F. Schneider. Coordinated multi-robot exploration. *IEEE Transactions on Robotics*, 21(3):376–386, 2005.

[10] J. L. Burke, R. Murphy, M. D. Coovert, and D. L. Riddle. Moonlight in miami: Field study of human-robot interaction in the context of an urban search and rescue disaster response training exercise. *Human-Computer Interaction*, 19(1-2):85–116, 2004.

[11] J. Casper and R. R. Murphy. Human-robot interactions during the robot-assisted urban search and rescue response at the world trade center. *Trans. Sys. Man Cyber. Part B*, 33(3):367–385, 2003.

[12] Eric C. Chi and Tamara G. Kolda. On tensors, sparsity, and nonnegative factorizations. *SIAM Journal on Matrix Analysis and Applications*, 33(4):1272–1299, 2012.

[13] H. Choi, L. Brunet, and J.P. How. Consensus-based decentralized auctions for robust task allocation. *IEEE Transactions on Robotics*, 25(4):912–926, 2009.

[14] Gerald Seet Choon Yue Wong and Siang Kok Sim. Multiple-robot systems for usar: Key design attributes and deployment issues. *International Journal of Advanced Robotic Systems,*, 8(1):85–101, 2011.

[15] A. Christensen, R. OGrady, M. Birattari, and M. Dorigo. Fault detection in autonomous robots based on fault injection and learning. *Autonomous Robots*, 24(1):49–67, 2008.

[16] R. Cipolleschi, M. Giusto, A. Q. Li, and F. Amigoni. Semantically-informed coordinated multirobot exploration of relevant areas in search and rescue settings. In *European Conference on Mobile Robots*, pages 216–221, 2013.

[17] Complexible. Stardog: Enterprise graph database. `http://stardog.com`, 2015.

[18] M. S. Couceiro, R. P. Rocha, and N.M.F. Ferreira. A novel multi-robot exploration approach based on particle swarm optimization algorithms. In *IEEE International Symposium on Safety, Security, and Rescue Robotics*, pages 327–332, 2011.

[19] Mattew Crosby. Adp an agent decomposition planner. `http://homepages.inf.ed.ac.uk/mcrosby1/`, 2015.

[20] J. De la Cruz, C. D. Paternina-Arboleda, V. Cantillo, and J. R. Montoya-Torres. A two-pheromone trail ant colony systemtabu search approach for the heterogeneous vehicle routing problem with time windows and multiple products. *Journal of Heuristics*, 19(2):233–252, 2013.

[21] A. Dietrich, S. Zug, S. Mohammad, and J. Kaiser. Distributed management and representation of data and context in robotic applications. In *IROS*, pages 1133–1140, 2014.

[22] Kais Dukes. Semantic annotation of robotic spatial commands. In *Language and Technology Conference LTC*, 2013.

[23] R.C. Farias, P. Comon, and R. Redon. Data mining by nonnegative tensor approximation. In *IEEE International Workshop on Machine Learning for Signal Processing*, pages 1–6, 2014.

[24] The Apache Software Foundation. Apache cassandra database. `http://cassandra.apache.org`, 2015.

[25] Michael Genesereth and Richard Fikes. Knowledge interchange format version 3.0 reference manual. Technical report, Stanford Logic Group Report, Stanford University, 1992.

[26] B. Grocholsky, R. Swaminathan, J. Keller, V. Kumar, and G. Pappas. Information driven coordinated air-ground proactive sensing. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, pages 2211–2216, 2005.

[27] A. Howard, L. E. Parker, and G. S. Sukhatme. Experiments with a large heterogeneous mobile robot team: Exploration, mapping, deployment and detection. *Int. J. Rob. Res.*, 25(5-6):431–447, 2006.

[28] M. Hsieh, A. Cowley, J. Keller, L. Chaimowicz, B. Grocholsky, V. Kumar, C. Taylor, Y. Endo, R. Arkin, B. Jung, D. Wolf, G. Sukhatme, and D. MacKenzie. Adaptive teams of autonomous aerial and ground robots for situational awareness. *Journal of Field Robotics*, 24(11-12):991–1014, 2007.

[29] M.E. Khatir and E.J. Davidson. Decentralized control of a large platoon of vehicles operating on a plane with steering dynamics. In *American Control Conference*, pages 2159–2165, 2005.

[30] G. Ayorkor Korsah, Anthony Stentz, and M. Bernardine Dias. A comprehensive taxonomy for multi-robot task allocation. *Int. J. Rob. Res.*, 32(12):1495–1512, 2013.

[31] Gregory Kuhlmann, Peter Stone, Raymond J. Mooney, and Jude W. Shavlik. Guiding a reinforcement learner with natural language advice: Initial results in robocup soccer. In *The AAAI-2004 Workshop on Supervisory Control of Learning and Adaptive Systems*, 2004.

[32] M. Lewis, Huadong Wang, Shih-Yi Chien, P. Scerri, P. Velagapudi, K. Sycara, and B. Kane. Teams organization and performance in multi-human/multi-robot teams. In *IEEE International Conference on Systems Man and Cybernetics*, pages 1617–1623, 2010.

[33] C.S. Lim, R. Mamat, and T. Braunl. Market-based approach for multi-team robot cooperation. In *4th International Conference on Autonomous Robots and Agents*, pages 62–67, 2009.

[34] Gi Hyun Lim, Il Hong Suh, and Hyowon Suh. Ontology-based unified robot knowledge for service robots in indoor environments. *IEEE Transactions on Systems, Man and Cybernetics, Part A: Systems and Humans*, 41(3):492–509, 2011.

[35] M. Lindner, M. Kalech, and G. A. Kaminka. A representation for coordination fault detection in large-scale multi-agent systems. *Annals of Mathematics and Artificial Intelligence*, 56(2):153–186, 2009.

[36] Yuanyuan Liu and Fanhua Shang. An efficient matrix factorization method for tensor completion. *IEEE Signal Processing Letters*, 20(4):307–310, 2013.

[37] R. MacGregor. Retrospective on Loom. Technical report, University of Southern California, Information Sciences Institute, Los Angeles, 1999.

[38] J. Manathara, P. B. Sujit, and R. Beard. Multiple uav coalitions for a search and prosecute mission. *Journal of Intelligent and Robotic Systems*, 62(1):125–158, 2011.

[39] A. Marino, L. Parker, G. Antonelli, and F. Caccavale. A decentralized architecture for multi-robot systems based on the null-space-behavioral control with application to multi-robot border patrolling. *Journal of Intelligent and Robotic Systems*, 71(3-4):423–444, 2013.

[40] M. Marneffe, B. Maccartney, and C. Manning. Generating typed dependency parses from phrase structure parses. In *Proceedings of the Fifth International Conference on Language Resources and Evaluation LREC*, May 2006.

[41] D. Milutinovic and P. Lima. Modeling and optimal centralized control of a large-size robotic population. *IEEE Transactions on Robotics*, 22(6):1280–1285, 2006.

[42] R. M. Murray. Recent research in cooperative control of multivehicle systems. *Journal of Dynamic Systems, Measurement, and Control*, 129(5):571–583, 2007.

[43] N. S. Nathan Srebro Nati and T. Jaakkola. Weighted low-rank approximations. In *ICML*, pages 720–727. AAAI Press, 2003.

[44] P. D. O'Brien and R. C. Nicol. Fipa — towards a standard for software agents. *BT Technology Journal*, 16(3):51–59, 1998.

[45] A. Ollero, S. Lacroix, L. Merino, Jeremi Gancet, J. Wiklund, V. Remuss, I.V. Perez, L.G. Gutierrez, D.X. Viegas, M.A.G. Benitez, A. Mallet, R. Alami, R. Chatila, G. Hommel, F.J.C. Lechuga, B.C. Arrue, J. Ferruz, J.R. Martinez-De Dios, and F. Caballero. Multiple eyes in the skies: architecture and perception issues in the comets unmanned air vehicles project. *IEEE Robotics Automation Magazine*, 12(2):46–57, 2005.

[46] L. E. Parker. Lifelong adaptation in heterogeneous multi-robot teams: Response to continual variation in individual robot performance. *Auton. Robots*, 8(3):239–267, 2000.

[47] L. E. Parker. Multiple mobile robot systems. In Bruno Siciliano and Oussama Khatib, editors, *Springer Handbook of Robotics*, pages 921–941. Springer Berlin Heidelberg, 2008.

[48] Lynne E Parker. Distributed intelligence: Overview of the field and its application in multi-robot systems. *Journal of Physical Agents*, 2(1):5–14, 2008.

[49] AH Phan and A Cichocki. Multi-way nonnegative tensor factorization using fast hierarchical alternating least squares algorithm HALS. In *Proc. of The 2008 International Symposium on Nonlinear Theory and its Applications*, 2008.

[50] Fiora Pirri. The well-designed logical robot: Learning and experience from observations to the situation calculus. *Artificial Intelligence*, 175(1):378 – 415, 2011. John McCarthy's Legacy.

[51] Jason D. M. Rennie, Lawrence Shih, Jaime Teevan, and David R. Karger. Tackling the poor assumptions of naive bayes text classifiers. In *In Proceedings of the Twentieth International Conference on Machine Learning*, pages 616–623, 2003.

[52] Craig Schlenoff and Elena Messina. A robot ontology for urban search and rescue. In *Proceedings of the 2005 ACM Workshop on Research in Knowledge Representation for Autonomous Systems*, KRAS '05, pages 27–34, New York, NY, USA, 2005. ACM.

[53] J.N. Schwertfeger and O.C. Jenkins. Multi-robot belief propagation for distributed robot allocation. In *International Conference on Development and Learning*, pages 193–198, 2007.

[54] Fabrizio Sebastiani. Machine learning in automated text categorization. *ACM Comput. Surv.*, 34(1):1–47, 2002.

[55] G. Simzan, A. Akbarimajd, and M. Khosravani. A market based distributed cooperation mechanism in a multi-robot transportation problem. In *11th International Conference on Intelligent Systems Design and Applications*, pages 1–5, 2011.

[56] A. Stroupe, T. Huntsberger, A. Okon, H. Aghazarian, and M. Robinson. Behavior-based multi-robot collaboration for autonomous construction tasks. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 1495–1500, 2005.

[57] Il Hong Suh, Gi Hyun Lim, Wonil Hwang, Hyowon Suh, Jung-Hwa Choi, and Young-Tack Park. Ontology-based multi-layered robot knowledge framework (omrkf) for robot intelligence. In *IROS*, pages 429–436, 2007.

[58] P B Sujit and D Ghose. Negotiation schemes for multi-agent cooperative search. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 223(6):791–813, 2009.

[59] Ilya Sutskever, Joshua B. Tenenbaum, and Ruslan R Salakhutdinov. Modelling relational data using bayesian clustered tensor factorization. In Y. Bengio, D. Schuurmans, J.D. Lafferty, C.K.I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 1821–1828. Curran Associates, Inc., 2009.

[60] P. Symeonidis, A. Nanopoulos, and Y. Manolopoulos. A unified framework for providing recommendations in social tagging systems based on ternary semantic analysis. *IEEE Transactions on Knowledge and Data Engineering*, 22(2):179–192, 2010.

[61] S. Tadokoro. Earthquake disaster and expectation for robotics. In S. Tadokoro, editor, *Rescue Robotics*, pages 1–16. Springer London, 2009.

[62] Jaime Teevan and David R. Karger. Empirical development of an exponential probabilistic model for text retrieval: Using textual analysis to build a better model. In *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Informaion Retrieval*, pages 18–25, New York, NY, USA, 2003. ACM.

[63] Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew R. Walter, Ashish G. Banerjee, Seth Teller, and Nicholas Roy. Understanding natural language commands for robotic navigation and mobile manipulation. In *Proceedings of the National Conference on Artificial Intelligence AAAI*, pages 1507–1514, 2011.

[64] M. Tenorth, A.C. Perzylo, R. Lafrenz, and M. Beetz. Representation and exchange of knowledge about actions, objects, and environments in the roboearth framework. *IEEE Transactions on Automation Science and Engineering*, 10(3):643–651, 2013.

[65] Alejandro Torreño, Eva Onaindia, and Óscar Sapena. Fmap: distributed cooperative multi-agent planning. *Applied Intelligence*, 41(2):606–626, 2014.

[66] Jan Tozicka, Jan Jakubuv, Karel Durkota, Antonín Komenda, and Michal Pechoucek. Multiagent planning supported by plan diversity metrics and landmark actions. In *ICAART*, pages 178–189, 2014.

[67] I. Vasilescu, K. Kotay, D. Rus, M. Dunbabin, and P. Corke. Data collection, storage, and retrieval with an underwater sensor network. In *Proceedings of the 3rd International Conference on Embedded Networked Sensor Systems*, SenSys '05, pages 154–165, New York, NY, USA, 2005. ACM.

[68] P. Velagapudi, P. Scerri, K. Sycara, Huadong Wang, M. Lewis, and Jijun Wang. Scaling effects in multi-robot control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2121–2126, 2008.

[69] L. Vig and JulieA. Adams. Coalition formation: From software agents to robots. *Journal of Intelligent and Robotic Systems*, 50(1):85–118, 2007.

[70] F. Zhang, W. Chen, and Y. Xi. Improving collaboration through fusion of bid information for market-based multi-robot exploration. In *IEEE International Conference on Robotics and Automation*, pages 1157–1162, 2005.

[71] T. Zhang and H. Ueno. A frame-based knowledge model for heterogeneous multi-robot system. *IEEE Transactions on Electronics, Information and Systems*, 125(6):846–855, 2005.

[72] Y. Zhang and L.E. Parker. Iq-asymtre: Forming executable coalitions for tightly coupled multirobot tasks. *IEEE Transactions on Robotics*, 29(2):400–416, 2013.

[73] H. elikkanat and E. ahin. Steering self-organized robot flocks through externally guided individuals. *Neural Computing and Applications*, 19(6):849–865, 2010.

Evaluation of Multi-Agent Planning frameworks under distributed privacy-preserving

Mario Gianni, Manuel A. Ruiz Garcia and Fiora Pirri
ALCOR, Vision, Perception and Learning Robotics Laboratory
DIAG, University of Rome 'La Sapienza', Italy

Abstract

In this report we describe the evaluation of current state-of-the-art Multi-Agent planning frameworks focusing on privacy-preserving. In these frameworks, agents can decide to share public capabilities while keeping private processes and information that support these capabilities. Several instances of well-known cooperative multi-agent problems have been considered for this evaluations. Plan quality, planning time, ratio of problem solved and dimension of public information have been used for measuring performance of the planners under consideration. Results and discussion are also reported, in particular, under the context of the EU FP7 ICT Project TRADR, grant no. 609763.

1 Introduction

In order to reason about both individual and joint actions a team of cooperative robots requires knowledge about the domain as well as knowledge about the state of each member of the team. Very often this knowledge is not available for several reasons. Domains in which real robots have to operate are usually partially observable. The communication channel supporting information exchange is unreliable.

Communication unreliability is in particular a crucial issue in TRADR system. During past end-users evaluations, we realized that, although network infrastructure has been boosted (WP6), it was not able to support the amount of information (e.g., point cloud, images) exchanged on the channel, thus causing freezing as well as loss of situation awareness among robots and operators.

This in-field experience highlighted that a crucial aspect which has to be accounted, before developing a planning algorithm for multiple cooperative robots, is the analysis of what data are really needed in order to reason about actions and of what information is not required to be necessary exchanged among the members of a team of robots. Preliminary dealing with this aspect it would avoid an overloading of the communication channel due to exchanged redundant information, thus leading to a more careful use of the channel bandwidth.

Recently, at the Competition of Distributed and Multi-Agent Planners (CoDMAP) [20], several algorithms have been presented accounting for what information has to be kept private as well as what data have to be conversely shared for cooperative agent planning [5, 12, 14]. In these algorithms, agents supply a public interface only and, through a distributed planning process, generate plans that achieves the desired goals without being required to share a complete model of their actions and local state with other agents.

These planners takes as input instances of multi-agent planning problems specified in the Multi-Agent STRIPS formalism [6]. MA-STRIPS is well-suited to model problems where the agents agree to cooperate to achieve a joint goal but do require to share some of their private information. Indeed, in MA-STRIPS each fact is classified either as public or as internal. A fact is internal for agent when it is not public but mentioned by some action of the agent. A fact is relevant for an agent when it is either public or internal for the agent. MA-STRIPS further extends this classification of facts to actions. An action is public when it has a public effect, otherwise it is internal. An action is relevant for an agent when it is either public or owned by the agent. Therefore, in multi-agent planning modeled as MA-STRIPS, agents can either plan only with their own actions and facts and inform the other agents about public achieved facts, or can also use other agents public actions provided that the actions are stripped of the private facts in preconditions and effects. Thus agents plan actions for other agents and then coordinate the plans [16].

In order to perform an analysis of the feasibility of planners accounting for both private and public information to TRADR domain and scenarios, we have taken four frameworks, which participated to the CoDMAP competition and we have evaluated their performance on six (out of nine) different MAP domain problems, taken form [10]. These four frameworks are

- Multi-Agent Planning by plan Reuse (MAPR) [4];
- Centralized Multi-Agent Planning (CMAP) [4];
- Agent Decomposition Planner (ADP) [8];

- MAPlan [9];
- Forward-Chaining Multi-Agent Planning (FMAP) [15];

In the following we provide a brief description of the main features of these planners as well as a brief description of the considered planning domains with respect to TRADR scenarios. We also introduce evaluation metrics, similar to those used in the competition and finally we report the results obtained according to these metrics.

2 Planners description

MAPR algorithm [2] starts by assigning all common goals between agents. Upon the completion of the assignment, the algorithm solves (iteratively) each agent problem. Once an agent finds a solution, the private component of the solution is obfuscated and communicated to the next agent, such that successive agent must solve its own problem augmented with both, the obfuscated private components and public ones, of all previous solutions. Under this perspective, MAPR threats multi-agent planning as plan reuse. Conversely, CMAP [3] lets each agent obfuscates its own private information, that will be sent to a centralized planning agent. This centralized agent performs a centralized planning step with all the obfuscated information provided by the set of agents.

ADP is a complete, non-optimal centralised planning algorithm that attempts to compute and utilise agent decompositions to improve planning time [7]. ADP system comprises two main components: a decomposition phase and a heuristic calculation. The decomposition phase is responsible to find a feasible multi-agent decomposition of the planning problem, such that agents can only influence themselves or the environment. Consequently, actions affecting the internal states of multiple agents are not allowed. Once a multi-agent decomposition is available, a greedy best first search is used to find plans faster than the single-agent approach. The greedy search is guided by a multi-agent heuristic calculation, based on the simple idea of generating planning graphs only for a single agent sub-problem at a time.

MAPlan planner extends the multi-agent A* for parallel and distributed systems (MAD-A*) [11]. It is based on both multi-threaded and distributed state space heuristic search. In MAPlan, each operator is assigned to a single agent and each agent expands the state space of its owned subset of operators. All the information regarding public operators is always shared to the rest of agents. This algorithm implements four different heuristics: LM-cut [1], distributed LM-cut [17], distributed Fast-Forward (FF) [19] and distributed FF [18]

The fully distributed multi-agent planner FMAP [14] is based on a complete and suboptimal distributed A*, that iteratively search a multi-agent tree. Each node of the multi-agent search tree is built with the contributions of one or more agents. The iterative procedure is based on a democratic leadership, assigning a coordinator role to a different single agent at each iteration. The heuristic guiding the searching procedure is based on the concepts of Domain Transition Graph and frontier state, optimized to evaluate plans in distributed environments.

3 Domains description

The experimental setup is composed by six different MAP domains taken from [10]. The remaining domains in [10] have not been considered due to missing of private information filtering. Therefore, evaluation with respect to TRADR scenarios is not relevant. Two domains (*Rovers* and *Satellite*) are loosely-coupled, that is, where agents have the same planning capabilities and thus, each task goal can be solved by any single agent without cooperating with the other. The other four (*Depots*, *Elevators*, *Logistics* and *Woodworking*) are tightly-coupled, meaning that solving a task requires interactions or commitments among agents. These domains have been considered for evaluation due to similarities with TRADR system. Indeed, multi-robot cooperative planning in TRADR can be loosely and tightly due to the deployment of homogeneous (UGV-UGV) and heterogeneous (UGV-UAV) robots. We now briefly describe each domain.

Rovers In this domain agent is assigned each rover (ranging from 1 up to 8). Rovers must collect samples of soil and rock and only interact between each other when a sample is collected, since such sample is no longer available for the rest of agents. Only the information regarding the collected samples is public.

Satellite As in the *Rovers* case, each satellite (ranging from 1 up to 12) correspond to an agent. Since each satellite can reach independently (without cooperation) a subset of the task goals, the resulting MAP tasks are almost decoupled. Location, orientation and instruments of the satellites are private to the agents and only the collected images by the satellites are defined as public.

Depots This domain introduces two specialized agents, depots and trucks, that must cooperate in order to solve most of the MAP task goals. Therefore, goals are tightly-coupled with many different dependencies among agents. Agents range starts from 5 up to 12. Only the location of packages and trucks are publicly available.

Elevators Agents (ranging from 3 up to 5) can be either a slow-elevator or a fast-elevator. Regardless the fact that the operators in this STRIPS domain are basically the same for both types of agents, elevator agents are still specialized since they can access (or not) different floors depending on its type, forcing cooperation to complete some task goals. The locations of different passengers are shared between agents.

Logistics In this domain, specialized agents (ranging from 3 up to 10) can represent either air-planes or trucks. Package delivery may require several forms of cooperation between the specialized agents. The location of packages is defined as public.

Woodworking Four different types of specialized agents are defined within this domain (planner, saw, grinder and varnisher) representing machines in a pipelined production chain. All task of the domain include four agents one of each type. Information regarding wood pieces is public since agent cooperation relies on it.

4 Evaluation metrics

Evaluation of the planners has been performed over 20 instances of the domain problems under consideration. For each domain and for each planner we report the ratio of solved instances, the quality Q of the generated plan and, finally, the planning time T. Both quality and time are measured according to the International Planning Competition (IPC) scoring. In particular, the planning quality Q is equal to the sum of all ratios Q_i^*/Q_i, i=1,…,20, with Q_i the plan cost of the problem i and Q_i^* the cost of the best plan found by any of the planners for the problem i. Plan cost Q_i is defined with respect to action cost. In all the instances of the problems unit-cost of action has been considered. On the other hand, for each problem i the score associated to the planning time is given by $T_i = \dfrac{1}{1+\log_{10}(T_i/T_i^*)}$. Here, T_i^* the minimum time required by any planner to solve the problem. A zero time score is assigned when a planner does not solve the problem. For each domain and for each planner we also report a measure D accounting for the size in bytes of the amount of information exchanged, that is public, for generating a plan. This measure is computed as the ratio between the size in bytes of the public information needed for computing the worst plan found by any planner and the average size of a planning problem.

5 Results and Discussion

Both Table 1 and Table 2 report for each domain and for each planner the number of solved instances, the quality Q of the generated plan and the planning time T. In particular, for CMAP and for MAPlan two variants and three variants of the algorithms have been evaluated, respectively. CMAP-t is the variant of CMAP algorithm endowed with a subset goal assignment strategy and LAMA-UNIT-COST as the base planner. CMAP-q uses as a base planner LAMA-2011 [13]. Variants of MAPlan algorithm differ from the underlying heuristic strategy (see Section 2). On the other hand, Table 3 reports the evaluation of the planners under consideration, with respect to the measure D of the amount of information exchanged for plan generation, as defined in Section 4

Domain	Problems	CMAP-t				CMAP-q				MAPR-p				MH-FMAP			
		Solved	Ratio	Q	T	Solved	Ratio	Q	T	Solved	Ratio	Q	T	Solved	Ratio	Q	T
Depots	20	17	0.85	12.60	15.95	17	0.85	17.00	6.20	0	0.00	0.00	0.00	2	0.10	1.88	0.76
Elevators	20	20	1.00	11.91	17.08	18	0.90	17.47	4.97	20	1.00	15.25	17.71	9	0.45	5.12	2.68
Logistics	20	20	1.00	18.81	17.51	19	0.95	18.38	5.39	20	1.00	19.98	14.43	4	0.20	3.88	1.61
Rovers	20	20	1.00	18.72	17.38	20	1.00	19.99	5.26	19	0.95	14.84	16.49	7	0.35	6.33	2.23
Satellites	20	20	1.00	19.22	17.37	20	1.00	19.75	5.17	19	0.95	16.43	16.99	17	0.85	16.61	6.09
Woodwork	20	16	0.80	14.98	14.34	15	0.75	13.71	5.10	0	0.00	0.00	0.00	8	0.40	7.06	3.08

Table 1: Comparison of MAP planning algorithms CMAP-t, CMAP-q, MAPR-p and MH-FMAP

The reported results highlight that CMAP-q turns out to be the planner algorithm with the best performance. Indeed, it optimizes both the quality of the generated plan and problem coverage. With respect to the other algorithms, it also requires less information to be exchanged in order to achieve the planning task. The main strength of this algorithm is the strategy used for goal assignment. According to this strategy, for each public goal, the planner computes the relaxed plan. Moreover, instead of just computing the cost, it computes the subset of agents that appear in the relaxed plan for that goal, namely, those are agents that could potentially be needed to achieve the goal. Another feature of this algorithm is a planner-independent strategy for plan

Domain	Problems	MAPLan/LM-Cut				MAPlan/MA-LM-Cut				MAPlan/FF+DTG				ADP			
		Solved	Ratio	Q	T	Solved	Ratio	Q	T	Solved	Ratio	Q	T	Solved	Ratio	Q	T
Depots	20	5	0.25	5.00	2.63	4	0.20	4.00	1.63	13	0.65	7.72	8.38	13	0.65	7.48	9.71
Elevators	20	2	0.10	2.00	1.01	2	0.10	2.00	0.72	12	0.60	5.15	5.60	20	1.00	12.99	16.41
Logistics	20	6	0.30	6.00	2.81	8	0.40	8.00	3.36	18	0.90	11.23	10.92	20	1.00	18.98	16.13
Rovers	20	1	0.05	1.00	0.38	1	0.05	1.00	0.36	20	1.00	17.88	14.82	20	1.00	18.31	13.89
Satellites	20	3	0.15	3.00	1.71	7	0.35	7.00	2.85	19	0.95	17.64	11.30	20	1.00	17.00	15.03
Woodwork	20	5	0.25	5.00	2.79	5	0.25	5.00	3.31	15	0.75	13.23	12.15	20	1.00	18.96	16.85

Table 2: Comparison of MAP planning algorithms MAPLan/LM-Cut, MAPlan/MA-LM-Cut, MAPlan/FF+DTG and ADP.

Domain	CMAP-t	CMAP-q	MAPR-p	MH-FMAP	MAPLan/LM-Cut	MAPlan/MA-LM-Cut	MAPlan/FF+DTG	ADP
	D	D	D	D	D	D	D	D
Depots	0.0620	0.0620	-	0.5270	0.2108	0.2635	0.0811	0.0811
Elevators	0.0439	0.0487	0.0439	0.0975	0.4386	0.4386	0.0731	0.0439
Logistics	0.1058	0.1114	0.1058	0.5290	0.3527	0.2645	0.1176	0.1058
Rovers	0.0072	0.0072	0.0075	0.0205	0.1432	0.1432	0.0072	0.0072
Satellites	0.0474	0.0474	0.0499	0.0558	0.3162	0.1355	0.0499	0.0474
Woodwork	0.0234	0.0250	-	0.0469	0.0750	0.0750	0.0250	0.0188

Table 3: Comparison of MAP planning algorithms with respect to the dimension in bytes of the public information needed for computing the worst plan.

parallelization. Indeed, this planner is endowed with a suboptimal procedure which generates a partially-ordered plan from a totally-ordered one. Then, a potential parallelization of the sequential plan is extracted from the partially-ordered plan. This procedure turns out to be a useful criteria for optimizing make-span. This is a crucial aspect when planning has to take into account on-line scheduling of the tasks. Despite its performance, CMAP-q is centralized. A planning paradigm which is based on a centralized agent for dispatching tasks to the other agents might be a bottleneck in TRADR system. Among the distributed planning algorithms that we have evaluated, on the basis of the results in Table 1, 2 and 3, a prominent alternative might be the variant of MAPlan with distributed Fast-Forward heuristic, based on Domain Transition Graphs (DTGs) [18]. This heuristic is based on an exploration of the local DTGs of the agents, constructed from projected operators. The possibly unknown preconditions and effects of projected operators are recorded into the DTGs with a special symbol. Once that symbol is reached during extraction of the relaxed plan, a distributed recursion is executed and the cost of the relaxed plan is computed with the support of the owners of the projected operators. Moreover, partial plans for each fact can be cached and later reused without re-computation of the full estimate. Together with distributed planning, both plan caching and reuse are important features that have to be accounted in TRADR system, especially for developing persistence in multi-robot collaboration.

Acknowledgment

This work reports research partially funded by the EU FP7 ICT Project TRADR, grant no. 609763.

References

[1] Blai Bonet and Malte Helmert. Strengthening landmark heuristics via hitting sets. In *ECAI*, pages 329–334, 2010.

[2] Daniel Borrajo. Multi-agent planning by plan reuse. In *AAMAS*, pages 1141–1142, 2013.

[3] Daniel Borrajo. Plan sharing for multi-agent planning. In *ICAPS*, pages 57–65, 2013.

[4] Daniel Borrajo and Susana Fernandez. Multi-agent planning by plan reuse. http://www.plg.inf.uc3m.es/resources.php, 2015.

[5] Ronen Brafman. A privacy preserving algorithm for multi-agent planning and search. In *IJCAI*, 2015.

[6] Ronen Brafman and Carmel Domshlak. From one to many: planning for loosely coupled multi-agent systems. In *ICAPS*, pages 28–35, 2008.

[7] Matt Crosby, Michael Rovatsos, and Ronald P. A. Petrick. Automated agent decomposition for classical planning. In *ICAPS*, 2013.

[8] Mattew Crosby. Adp an agent decomposition planner. http://homepages.inf.ed.ac.uk/mcrosby1/, 2015.

[9] Daniel Fišer, Michal Štolba, and Antonín Komenda. Maplan. https://github.com/danfis?tab=repositories, 2015.

[10] Derek Long and Maria Fox. The 3rd international planning competition: results and analysis. *JAIR*, 20:1–59, 2003.

[11] Raz Nissim and Ronen I. Brafman. Multi-agent a* for parallel and distributed systems. In *AAMAS*, pages 1265–1266, 2012.

[12] Raz Nissim and Ronen I. Brafman. Distributed heuristic forward search for multi-agent planning. *JAIR*, 51:293–332, 2014.

[13] Silvia Richter and Matthias Westphal. The lama planner: Guiding cost-based anytime planning with landmarks. *J. Artif. Int. Res.*, 39(1):127–177, 2010.

[14] Alejandro Torreño, Eva Onaindia, and Óscar Sapena. Fmap: distributed cooperative multi-agent planning. *Applied Intelligence*, 41(2):606–626, 2014.

[15] Alejandro Torreño, Óscar Sapena, and Eva Onaindia. Fmap: forward-chaining multi-agent planning system. http://users.dsic.upv.es/grupos/grps/tools/map/fmap.html, 2015.

[16] Jan Tozicka, Jan Jakubuv, Karel Durkota, Antonín Komenda, and Michal Pechoucek. Multiagent planning supported by plan diversity metrics and landmark actions. In *ICAART*, pages 178–189, 2014.

[17] Michal Štolba, Daniel Fišer, and Antonín Komenda. Admissible landmark heuristic for multi-agent planning. In *ICAPS*, pages 211–219, 2015.

[18] Michal Štolba, Daniel Fišer, and Antonín Komenda. Comparison of rpg-based ff and dtg-based ff disrtibuted heuristics. In *Proceedings of the 3rd Workshop on DMAP*, pages 77–82, 2015.

[19] Michal Štolba and Antonín Komenda. Relaxation heuristics for multiagent planning. In *ICAPS*, 2014.

[20] Michal Štolba, Antonín Komenda, and Daniel Laszlo Kovacs. Competition of distributed and multi-agent planners (codmap) website. http://agents.fel.cvut.cz/codmap/, 2015.