

Adaptive Risk-Based Replanning for Human-Aware Multi-Robot Task Allocation with Local Perception

Zeynab Talebpour and Alcherio Martinoli

Abstract—In this paper, we propose an adaptive risk-based replanning strategy in the context of multi-robot task allocation for dealing with limitations of local perception and unpredicted human behavior. Our replanning method is based on the variations of social risk and human motion prediction uncertainty. The performance of our method is studied through an extensive suite of experiments of increasing complexity. Results obtained using both a high-fidelity simulator and real robots confirm that this strategy outperforms a non-adaptive replanning strategy in all cases with respect to the chosen social metrics. The overall performance of the team depends firstly on its replanning strategy, and secondly on the available information about the humans. Although an adaptive replanning strategy with global perception leads to the best performance, it is computationally expensive and infeasible in some real applications. Local perception shows comparable results as long as updates of relevant human poses affecting a task's risk are available within the execution time of that task. Conversely, the non-adaptive replanning strategy is shown to have degraded results with global perception as decisions in this case can be based on outdated information that lead to invalid plans.

Index Terms—Multi-Robot Systems, Social Human-Robot Interaction, Planning, Scheduling and Coordination.

I. INTRODUCTION

RESEARCH in the area of socially-aware navigation has received substantial attention in the recent years [1], [2]. Despite the numerous applications of socially-aware mobile robots as personal assistants at homes, robot tutors at schools, and service robots at hospitals and nursing homes, research in the human-aware navigation area focuses mainly on single robots and the problem of cooperative human-aware navigation for Multi-Robot Systems (MRS) is largely unexplored.

In this paper, we focus on a particular class of MRS coordination mechanisms commonly known as Multi-Robot Task Allocation (MRTA) [3] in social environments using a market-based approach [4]. In such environments the number of robots are often limited and the number of tasks are usually moderate. The main difficulty for MRTA in such highly dynamic and noisy environments is that plans are likely to change or to be rendered invalid, particularly, if the robots are planning for long periods of time. This is due to the highly stochastic nature of the problem, uncertainty inherent to human behavior, and limitations of robot perception. Additionally,

robots are required to perform in a socially acceptable manner in terms of navigation and interaction. This adds additional constraints to the planning problem.

Social MRTA takes humans into consideration in planning and coordination. Employing risk-based bids that have both predictive and social components, is one solution to this problem [5]. However, in real environments, global perception of the human poses is often not available and accessible to all robots. Hence, robots must rely on their local perception for decision making and planning instead. Moreover, there can be cases where the predictions are incorrect or not sufficiently accurate since humans can have sudden behavioral changes. Thus, the following questions need to be answered: i) how should a robot or a team of robots deal with the new or lost perceptual information that is inherent to local perception? ii) how should a robot react to unforeseen behavior changes of humans at the task planning level? We believe both of these problems should be addressed with a similar strategy.

In other words, the main MRTA challenge to be addressed is to endow a team of robots with an appropriate replanning strategy that allows for facing high uncertainties mainly generated by stochastic human-related information. In [6] such dilemma is summarized as follows: “Is it more beneficial to build a complex model that incorporates uncertainty, or is it enough to build less well-informed plans and replan as often as needed to quickly react to unexpected events?” We believe that a hybrid approach can also be taken for tackling this problem.

We employ a Hoplites-based MRTA framework [7] as the basis of our human-aware MRTA approach. The contributions of this paper include proposing an Adaptive Risk-Based Replanning (ARBR) strategy for handling i) new information and ii) unpredicted human behavior, leveraging the concept of risk-based bid estimation for human-aware coordination introduced in [5] and human motion uncertainty. ARBR enables the robots to modify their active plans by incorporating the new relevant updated information about humans in a distributed fashion. Furthermore, we perform an extensive suite of experiments in simulation and reality to evaluate the performance of the MRS employing this strategy. To the best of our knowledge, MRTA with limited perception and adaptive replanning in social environments has not been investigated in the literature.

II. RELATED WORK

MRTA algorithms vary in design and application [6], but their common objective is to find a mapping between robots in a team and a set of “tasks” that must be accomplished in order for the team’s “goal” to be completed. Among multiple approaches proposed for MRTA, we are mainly interested in distributed approaches that can be executed by a team of robots

Manuscript received: February, 24, 2019; Revised May, 17, 2019; Accepted June, 14, 2019. This paper was recommended for publication by Allison Okamura upon evaluation of the Associate Editor and Reviewers’ comments.

Zeynab Talebpour and Alcherio Martinoli are with the Distributed Intelligent Systems and Algorithms Laboratory (DISAL), School of Architecture, Civil and Environmental Engineering, École Polytechnique Fédérale de Lausanne (EPFL), 1015 Lausanne, Switzerland, firstname.lastname@epfl.ch. Videos accompanying this paper can be found at: <https://disal.epfl.ch/research/socialroboticsnavigation/>

without the explicit need for a centralized entity outside the team with perfect knowledge of the environment [8].

MRTA in social environments should ensure social acceptance as well as achieving good performances in terms of global team objectives such as the traveled distance, mission time, etc. In such environments, humans should be considered as social beings and not solely as dynamic obstacles in navigation. Respecting personal spaces, O-spaces and P-spaces [9] are the common social behaviors considered in the literature of human-aware navigation based on the concept of “proxemics” [10]. In this work, proxemics-based social costs encoded as costmaps similar to [11] along with a Fast Marching Method (FMM)-based path planner constitute the human-aware navigation components of the robots. FMM has been proven to be successful in real domestic spaces with high complexity [12]. There have previously been a number of research papers addressing social path planning using FMM [13].

One key element of dynamic social environments is change. A common solution for handling change in the literature of MRTA is replanning. A list of changes impacting the solution of MRTA can be found in [14]. The planning loop executed by most MRTA methods comprises planning-execution-replanning of tasks [6]. Replanning plays an important role in MRTA as it is responsible for handling these changes in the environment and maintaining some level of efficiency.

Typically, replanning occurs continuously or at predetermined points in time. It can also be triggered based on a set of events such as a robot accomplishing a task or arrival of a new task. However, not all changes or uncertainties in the environments can be captured efficiently using these strategies. A proactive replanning strategy that predicts the problems or opportunities in a construction assembly scenario using a centralized planner is proposed in [15]. Consensus Based Bundle Algorithm with Partial Replanning (CBBA-PR) is introduced in [16] for allocating new tasks that appear online during the solving of the task allocation problem. However, the focus of replanning in these works is not the stochastic changes in the cost of particular tasks or resolving the problem of plan change as a result of new information. Moreover, there is no human factor involved.

As social environments are very dynamic, the appropriate use of information updates is key in having a good performance. However, not every new update requires the robots to modify their plans and only a subset of these updates are relevant to the robot team. Having a replanning strategy that is activated upon every arrival of new information is very suboptimal. Replanning with a specified frequency is also prone to low performance: if the frequency is chosen to be too low with respect to the changing dynamics of the environment, the robots might not be able to react to the changes in a timely fashion; if the frequency is too high, this will result in a very resource-consuming operation.

III. PROBLEM DESCRIPTION AND APPROACH

In this section, we define the MRTA problem and explain the underlying market-based framework for robot coordination as well as the replanning strategy previously adopted. Next, we propose our adaptive risk-based replanning strategy.

A. Problem Description

Consider a team of n_r robots $\{r_j, j = 1, \dots, n_r\}$ and a set of n_t tasks $\{t_i, i = 1, \dots, n_t\}$, where each task is a location that a robot has to visit. The team of robots should decide how to efficiently subdivide these tasks using a local criterion in order to optimize a global performance. This global performance can be a function of time, distance traveled, etc. Tasks can be identified locally by the robots through on-board perception or can be broadcasted to all robots by an external source.

B. Social Multi-Robot Coordination

By means of a Hoplites-based coordination scheme [17] comprised of two main concurrent types of coordination, *passive coordination* and *active coordination*, the robots, which are self-interested agents in pursuit of individual profit, can evaluate each available task and decide whether to take it or sell/buy it to/from another robot in a distributed manner. Hoplites allows for coordinating *plans*, i.e., a sequence of tasks instead of only a single task.

In *passive coordination*, each robot chooses its most profitable plan and broadcasts it to other teammates without any attempt to modify their plans. This information is then used by other robots to reevaluate the expected profitability of their current plans, update and broadcast the changes. However, sometimes a robot's best plan can only be marginally profitable while a genuine team plan could result in a higher profit. This suggests that modifying the plans of the robot's teammates would be an interesting option to pursue. Concretely, this implies that the requesting robot asks its collaborators for compensation price quotes and persuades them to engage in a cooperative action. This process is ruled by a market-based approach and constitutes the *active coordination*. For details about our implementation of an Hoplites-based MRTA refer to [18].

The decision of switching to the active coordination mode is based on the evaluation of a *local balance function*. This function is strongly problem-dependent and can contribute to reaching the globally optimal solution at the team level if chosen correctly. The local balance functions constitute what is known as a bid in the context of the market. While many formulations can be used for encoding the local balance function of robots, operating in social environments calls for including a social factor in bids. Additionally, a problem-dependent *global balance function* is also required for team-level evaluations.

Perfect prediction of the future for humans is not possible because of the uncertainties inherent to uncontrolled social environments. Therefore, we propose an abstraction that can extract higher level information from perceptual data, by means of risks [5]. Risk forms the basis of our human-aware decision making and replanning, and is defined as the probability of occupation of an area with social costs by the robot. Risk is computed by means of expectation-based social costmaps [11] that incorporate the uncertainty in the human pose reported by a human trajectory predictor.

The local balance function of robot r_j , for each task t_i belonging to a plan P at time step k is defined in the following.

$$B_{r_j, P_k} = R_{r_j, P_k} - \left(C_{r_j, P_k} + \sum_{t_i \in P_k} f(\gamma_{r_j, t_i, k}) \right) \quad (1)$$

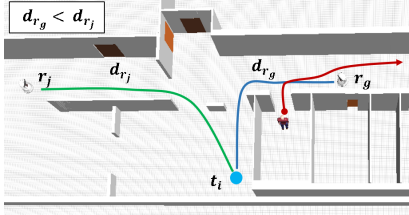


Fig. 1: A scenario where multiple active coordination attempts are required.

R is the revenue of the plan, C is a generic cost function that in this case is proportional to the length of the path planned by the FMM to the tasks in a given plan P . γ is the risk associated with t_i for r_j given the plan and $f(\cdot)$ is a user-defined monotonically increasing function used to aggregate the risk with revenue and cost terms.

The risk imposed by t_i for r_j consists of the risk at each waypoint w in the robot's trajectory W . By assigning costs on the basis of social costmaps to the risk formulation, social factors are further reinforced. Since the robot knows its planned path to t_i , it can compute the social risk at any w at time k_w by predicting where the humans ($h \in H$) will be at that time step. This is done by convolving a 2D Gaussian function N modeling the social costs, with p_h , the probability distribution of any human h being in the vicinity of that waypoint [5]. p_h is given by the human motion predictor and $[x_w, y_w]$ indicates the map coordinates of w .

$$\gamma_{r_j, t_i} = \sum_{w \in W_{r_j, t_i}} \sum_{h \in H} \int \int N(x_w - x, y_w - y) p_h(x, y) dx dy \quad (2)$$

C. Non-Adaptive Replanning

The strategy adopted by the robots in the non-adaptive Hoplite-based method is to replan when i) a task is accomplished and the robot is ready to take its next task, and ii) for verifying the validity of a stored plan when a robot is on its way towards a designated task. As long as an unfinished task t_i exists, robots perform replanning even if the remaining task is assigned to another robot r_j . If at any point in time replanning results in another robot r_g to be the best candidate for accomplishing t_i , active coordination is executed and t_i is delegated to the newly chosen robot r_g . This can be seen in Figure 1 where t_i is initially assigned to r_j at time step k as r_j plans first. However, as $d_{r_g} < d_{r_j}$ and $c_{t_i, r_g} < c_{t_i, r_j}$, robot r_g will request active coordination in its turn. We note that the human is initially static and starts moving after a while.

Regardless of the frequency at which replanning takes place, in the original Hoplites framework there is a constraint imposed by active coordination that prevents a robot from modifying its plan in some cases. In fact, active coordination can only be done *once* for a given task t_i between two robots r_j and r_g . When a robot accepts to participate in active coordination, it is bound by contract to do as promised. This is part of the Hoplites framework design to ensure that when the requesting robot r_g pays a compensation price to a robot r_j engaged in active coordination, an agreement is made based on which the desired task t_i will be assigned to the requesting robot r_g . In an environment with deterministic costs (*e.g.*, static obstacle only), this choice does not limit the robots. However, this limitation must be addressed for stochastic environments. In the context of our risk-based replanning, adopting this constraint can lead to a deadlock since when r_g senses the human motion, ARBR

is triggered and r_g will stop and cancel t_i . If r_g is stopped and knows r_j should take t_i but does not allow it, none of the robots will progress and thus a deadlock occurs for as long as $c_{t_i, r_j} < c_{t_i, r_g}$. We will address this problem by means of identifying cases in which active coordination must be permitted through risk monitoring. Our proposed risk-based replanning procedure will be working along side the replanning method previously available.

D. Adaptive Risk-Based Replanning

We would like to devise an adaptive replanning strategy that avoids replanning when unnecessary while being able to correctly identify when team plans should be revisited. In addition to active coordination as in Section III C, this strategy is realized by means of i) information sharing, ii) monitoring the social risks, and iii) risk-based rebidding when necessary.

1) *Information Sharing*: As most robotic systems only have limited and local perception in reality, they also have only a partial view of the environment. However, when multiple robots are distributed in different parts of a large social environment, they can provide important information to other team members about areas that may be out of their reach in terms of perception.

For estimating the stochastic component of costs in social MRTA, the main information to be communicated to team members is related to humans. For every perceived human h this information must include the human pose l_h . Other relevant information such as human velocity v_h , the interactions that the human is involved in I_h , etc., can all be reconstructed based on l_h with sufficiently fast perception updates. In our experiments, we implemented information sharing for a robot as sending pose information of directly perceived human targets to all team members and receiving pose information of human targets only perceived by other team members. Extending the perceptual domain of the team is not the only reason why information sharing is vital. If robots are taking decisions based on different assumptions to fulfill a collective goal, they must make sure there are no discrepancies in the information that forms the basis of their decisions. If such discrepancies exist, suboptimal decisions will be taken and specifically in our implementation of risk-based replanning, deadlocks can occur. As an example assume Figure 1 with local perception. Despite permitting multiple active coordinations, if r_g stops without communicating the increased social cost to r_j , both robots will remain idle. This leads to a deadlock for as long as $c_{t_i, r_j} < c_{t_i, r_g}$ and the human is only observed (and not communicated) by r_g .

2) *Risk Monitoring*: How do we decide if we should stop, step aside or continue walking when moving towards a destination in a crowded environment? We observe, predict and take an action. If something unexpected happens we are ready to adapt to the situation either by adjusting our path, waiting, or completely changing our route. What motivates us to modify our action at any point in time is the cost we estimate in an uncertain situation. In the context of MRTA, robots can take a similar approach by monitoring the social risk and the human motion uncertainty affecting every unfinished task in the environment locally.

We note that, for lost human targets who are no longer perceived by the robot, we perform motion prediction based on our human motion model to have a smooth transition in the social risk, since we need a realistic estimation of costs induced

by lost human targets despite not being able to perceive them. There is a limited time for lost target tracking after which the target is ignored since the uncertainty of prediction grows too large to be meaningful with lack of information updates.

3) *Risk-Based Rebidding*: As only a subset of information updates require the robots to revisit their plans, for identifying when a robot should change its current plan, the corresponding conditions must be defined. These conditions will be checked for triggering a rebidding command.

In general, a robot should reconsider its plan upon arrival of new information if risk estimation has a minimal level of uncertainty and i) the risk of accomplishing the currently active task is increasing, or ii) the risk of accomplishing another task has largely decreased making it more profitable than the currently active task. To formulate these conditions, we consider how the risk associated to a task can vary over time and how this change can affect the robot plan.

For a task t_i and a robot r_j , we denote the last time step that a bid has been placed for t_i prior to task allocation by k_{b,t_i} , the risk of t_i for r_j by γ_{t_i,r_j} , the gradient of risk indicating the rate of risk variation by $\dot{\gamma}_{t_i,r_j}$, the risk associated to t_i in the last bidding attempt by γ_{b,t_i} , and the cost of accomplishing t_i for r_j by c_{t_i,r_j} .

As the basis of risk monitoring is to react to change, we must ensure that risk variation is large enough to truly indicate a change. Additionally, with noisy perception and abrupt changes in the environment we must make sure that the variation in the risk trend is meaningful. Moreover, to avoid being too reactive to risk variations and repeatedly triggering rebidding, there must be a sufficiently large time window between rebidding attempts. These constraints constitute the first set of conditions written as follows:

$$|\gamma_{f,t_i} - \gamma_{b,t_i}| \geq \max(\alpha\gamma_{b,t_i}, \Gamma_{min}), \quad (k - k_{b,t_i}) \geq K \quad (3)$$

where γ_{f,t_i} is the filtered risk signal and K is a parameter indicating the minimum time interval between rebids. We implemented a median filter and average filter to remove outliers and small local variations. Γ_{min} is a minimum risk value and α is introduced to adapt the minimum risk value threshold in proportion to the risk magnitude.

Another key factor in risk-based rebidding is uncertainty. We consider moving humans to have a fairly high level of uncertainty associated to their trajectory prediction. However, if the human trajectory can be estimated with sufficient certainty due to lack of motion, then the variation of social risk would only be due to the increased social cost of the robot approaching a static human. In this case, γ_{t_i,r_j} is only considered for rebidding if a static human impacting c_{t_i,r_j} has been detected for the first time. We define σ_{t_i,r_j} as the social uncertainty associated to t_i for r_j . It indicates the maximum velocity among all human targets who impact the social risk γ_{t_i,r_j} . The minimum uncertainty level for ARBR is denoted as Σ . The second set of constraints are defined as follows:

$$\sigma_{t_i,r_j} > \Sigma, \quad \exists h \in D_k \wedge \gamma_{h,t_i,r_j} > 0 \quad (4)$$

where D_k is the set of newly perceived static humans for t_i at time step k and γ_{h,t_i,r_j} is the γ_{t_i,r_j} considering only h . The purpose of this condition is to include the risk of a static human only once in rebidding computations.

The active state of r_j and the assignment of t_i to r_j at time step k can be defined as follows respectively:

$$s_{r_j} \in \{1, 0\}, \quad s_{t_i,r_j} \in \{1, 0\} \quad (5)$$

Algorithm 1 Adaptive Risk-Based Replanning (ARBR) for robot r_j with a set of unfinished tasks T and a set of perceived humans H

```

1: procedure ARBR( $T, H$ )
2:   for  $t_i \in T$  do
3:     ▷ Compute the path to the location of  $t_i$  in an empty map
4:      $(W_{t_i}, K_{t_i}, d_{t_i}) \leftarrow \text{PathPlanning}(l_{r_j}, l_{t_i})$ 
5:     ▷ Compute the risk and uncertainty of human motion for  $t_i$ 
6:      $(\gamma_{t_i}, U_{H,t_i}) \leftarrow \text{RiskEstimation}(W_{t_i}, K_{t_i}, l_{t_i}, H)$ 
7:      $\Gamma_{t_i,k} \leftarrow \gamma_{t_i}$ 
8:      $c_{t_i,k} \leftarrow \text{ComputeCost}(\gamma_{t_i}, d_{t_i})$ 
9:     if  $B_{t_i} \neq \emptyset$  and  $t_a \neq \emptyset$  then
10:       $\Gamma_{f,t_i} \leftarrow \text{FilterRisk}(\Gamma_{t_i})$ 
11:      if RBRT( $\Gamma_{f,t_i}, U_{H,t_i}, c_{t_i,k}, t_i$ ) then
12:        Stop()
13:      ActiveCoordination( $t_i$ )
```

Algorithm 2 Risk-Based Rebid Triggering (RBRT) for robot r_j and task t_i given the filtered risk signal Γ_{f,t_i} , perceived human uncertainties U_{H,t_i} , and the task cost $c_{t_i,k}$ at time step k

```

1: procedure RBRT( $\Gamma_{f,t_i}, U_{H,t_i}, c_{t_i,k}, t_i$ )
2:    $f_v \leftarrow \text{False}, f_i \leftarrow \text{False}, f_a \leftarrow \text{False}$ 
3:    $f_c \leftarrow \text{False}, f_h \leftarrow \text{False}$ 
4:    $\gamma_{f,t_i} \leftarrow \Gamma_{f,t_i,k}$ 
5:    $\Gamma_{f,t_i} \leftarrow \text{ComputeRiskTrend}(\Gamma_{f,t_i})$ 
6:    $\dot{\gamma}_{f,t_i} \leftarrow \dot{\Gamma}_{f,t_i,-1}$  ▷ The latest risk variation rate
7:    $\sigma_{t_i} \leftarrow \max(U_{H,t_i})$ 
8:   ▷ Note any newly perceived static human that results in risk for  $t_i$ 
9:   if  $(\exists h \in D_k \wedge \gamma_{h,t_i} > 0)$  then
10:     $f_h \leftarrow \text{True}$  ▷ Human first encounter flag
11:   if  $|\gamma_{f,t_i} - \gamma_{b,t_i}| \geq \max(\alpha\gamma_{b,t_i}, \Gamma_{min}) \wedge (k - k_{b,t_i}) \geq K$  then
12:     $f_v \leftarrow \text{True}$  ▷ Risk variation flag
13:   else
14:     return False
15:   if  $\neg s_{t_i} \wedge (\dot{\gamma}_{f,t_i} < 0) \wedge (\sigma_{t_i} > \Sigma \vee f_h)$  then
16:     $f_i \leftarrow \text{True}$  ▷ Inactive task reconsideration flag
17:   if  $s_{t_i} \wedge (\dot{\gamma}_{f,t_i} > 0) \wedge (\sigma_{t_i} > \Sigma \vee f_h)$  then
18:     $f_a \leftarrow \text{True}$  ▷ Active task reconsideration flag
19:   if  $\neg s_{t_i}$  then
20:     $f_c \leftarrow (c_{t_i,k} < c_{t_a,k})$  ▷ Reduced cost flag
21:   else
22:     $f_c \leftarrow \text{True}$ 
23:   return  $f_v \wedge (f_i \vee f_a) \wedge f_c$  ▷ Final decision
```

If no task is allocated to a robot ($s_{r_j} = 0$) or there is no prior bid estimation ($B_{t_i,r_j} = \emptyset$) for t_i , there is no need for r_j to perform ARBR as there is no decision to be reconsidered. For an active robot ($s_{r_j} = 1$) however, the replanning decision for t_i depends on s_{t_i,r_j} .

In general, each robot tries to find a plan that minimizes c_{t_i,r_j} . For a given t_i this translates to minimizing γ_{t_i,r_j} . Therefore, an increasing risk ($\dot{\gamma}_{t_i,r_j} > 0$) is problematic if t_i is the currently active task ($s_{t_i,r_j} = 1$) and a decreasing risk ($\dot{\gamma}_{t_i,r_j} < 0$) is interesting if t_i is not assigned to the robot ($s_{t_i,r_j} = 0$). In this case, t_i will replace the currently active task only if it has a lower cost ($c_{t_i,r_j} < c_{t_a,r_j}$), where t_a is the currently active task. The last set of constraints considered for triggering rebidding can be written as follows:

$$s_{t_i,r_j} \wedge (\dot{\gamma}_{f,t_i} > 0), \quad \neg s_{t_i,r_j} \wedge (\dot{\gamma}_{f,t_i} < 0) \wedge (c_{t_i,r_j} < c_{t_a,r_j}) \quad (6)$$

The adaptive replanning and rebid triggering algorithms are detailed in Algorithm 1-2. We note that the subscripts indicating the robot are omitted for brevity as each algorithm is running locally on one robot.

IV. METRICS AND NAVIGATION

In this section, we will describe the evaluation metrics, navigation method, the robotic platform and the simulation

TABLE I: RBRT algorithmic parameters.

Parameter	α	Γ_{min}	K	Σ
Value	0.1	5	2	0.05

tool used for our experiments.

A. Metrics

For the global balance function concerning MRTA, the total traveled distance (M_1) and the mission time (M_2) are reported for all experiments. For evaluating the performance of the MRS in terms of social-awareness, the maximum accumulated social cost (M_3), the maximum time steps spent in areas associated with social costs (M_4) and the minimum distance to any human throughout the experiment (M_5) are reported among all robots.

B. Navigation

The navigation system is that of the MONarCH project [19], detailed in [12]. As input, it uses the pose estimates provided by a standard Adaptive Monte Carlo Localization (AMCL) self-localization system, given odometry, laser range finder readings, and a static map. The navigation system is based on the FMM for motion planning, together with a Dynamic Window Approach (DWA) algorithm for guidance and obstacle avoidance.

C. Robots

The robotic platform used in this work is called the MBot [20]. This is an omni-directional drive robot with an approximately round footprint of 0.65 m in diameter and a height of 0.98 m, endowed with two RGB-D cameras on the front with a conic FOV of 65° and 4 m of range. It is also equipped with two laser range finders, on both the front and the back, mounted 0.13 m above the ground, for providing 360° coverage with a range of 4 m.

D. Simulations

The use of high-fidelity simulators such as Webots is fundamental, especially when considering multi-robot systems. We have developed models of the environments (see Figure 1-4) for our experiments that enable simulations with similar environmental richness to the real world experiments, as shown in [18]. We rely on a pool of human trajectories with different walking speeds that can be (i) manually generated based on the knowledge of the designer (these trajectories can be very diverse as any sequence of way-points can be used to construct them), or (ii) automatically reproduced through the possibility of playing back real ROS bags with recorded human trajectories in a specific environment. In the risk calculation, a Kalman Filter is used for human trajectory prediction and the predicted human pose is a Gaussian distribution. Additionally, a constant velocity dynamics model is assumed for the human motion.

V. EXPERIMENTS

This section details the set of test cases used for our experiments. Each scenario has been repeated for ten runs. Robots are relying on their self-localization for computing the local balance functions. The evaluation metrics ($M_1 - M_5$) have been obtained from ground truth values provided by the simulation or a Motion Capture System (MCS) with

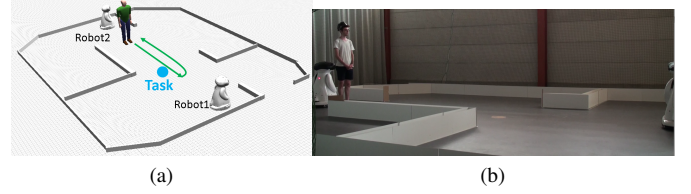


Fig. 2: Placement of the robots and the task in the arena for test case R-I. a) human walking trajectory, and b) snapshot of the initial state of the real robot experiment.



Fig. 3: Placement of the robots and the tasks for test case R-II. a) human walking trajectories, b) snapshot of the initial state of the real robot experiment.

millimetric tracking accuracy in real robot tests. The social metrics ($M_3 - M_5$) have been computed for the moving robots to avoid penalizing a static robot when a human decides to approach it. Throughout runs we have introduced randomness in human behavior by adding a random starting delay to the motion of each human. The algorithmic parameters of RBRT are shown in Table I. The local perception of the robots is currently emulated by the MCS in order to meet the specifications of the actual RGB-D camera mentioned above.

A. Test Case R-I: Multi-Robot Single Human Real Tests

This test case is designed to show how, despite having global perception, human behavior change can lead to a suboptimal plan for the robots. This test case consists of one task, two robots and one human shown in Figure 2. Four scenarios are considered in this test case, *Scenario NG*, non-adaptive replanning with global perception, *Scenario AG*, adaptive risk-based replanning with global perception, *Scenario NL*, non-adaptive replanning with local perception, and *Scenario AL*, adaptive risk-based replanning with local perception.

B. Test Case R-II: Multi-Robot Multi-Human Real Tests

In this test case, three tasks must be accomplished by a team of two robots in the presence of two walking humans that cannot be initially perceived by the robots. As a result, the initial risk-based bid estimations will not remain valid throughout the experiment. This test case highlights how information updates can be used to find better team plans. Figure 3 shows the simulated and real test environments of this test case. Four scenarios are considered in this test case. *Scenario NG*, non-adaptive replanning with global perception, *Scenario AG*, adaptive risk-based replanning with global perception, *Scenario NL*, non-adaptive replanning with local perception, and *Scenario AL*, adaptive risk-based replanning with local perception.

C. IPOL Simulations

In order to evaluate our method in more complex settings, an extensive suite of experiments has been conducted in an

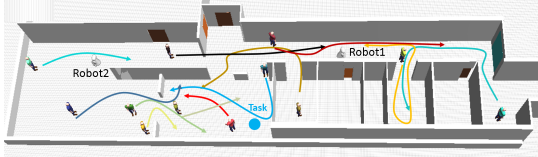


Fig. 4: The simulated arena of a sample test case with 12 moving humans. This snapshot shows the initial position of the robots and the trajectories of the humans.

environment that faithfully replicates the pediatric ward of an oncological hospital (IPOL)¹. The number of robots, humans, and tasks have been varied and different dynamics for the human motion have been considered in these experiments. Scenarios *NG*, *AG*, *NL* and *AL* have been tested for these settings similar to the previous test cases with different human motion dynamics.

VI. RESULTS

In this section the results of the test cases explained previously will be discussed. Before going to the details of the results, one emergent behavior observed during runs should be explained. Robots displayed a waiting behavior when confronted with increasing social costs for their currently active task. What drives this behavior for a robot r_j and an active task t_a is c_{t_a, r_j} :

$$c_{t_a, r_j} = \hat{c}_{t_a, r_j} + f(\gamma_{t_a, r_j}) \quad (7)$$

As r_j progresses towards t_a the deterministic part of the cost \hat{c}_{t_a, r_j} , is reduced and any increase in c_{t_a, r_j} would be due to the increase in γ_{t_a, r_j} . We note that $f(\cdot)$ is a monotonically increasing function. For $\dot{\gamma}_{t_a, r_j} > 0$ there can be cases where despite ARBR stopping the robot, r_j would still be the best candidate for t_a as $r_j = \text{argmin}_r(c_{t_a, r})$. Thus, t_a will be assigned to r_j again. Nonetheless, ARBR will stop r_j once again and the robot will be prevented from moving towards the assigned task. This results in a waiting behavior for as long as:

$$s_{r_j, t_a} = 1 \quad \wedge \quad \dot{\gamma}_{t_a, r_j} > 0 \quad \wedge \quad r_j = \text{argmin}_r(c_{t_a, r}) \quad (8)$$

The duration of the waiting period depends on how the environment changes. These waiting periods increase the total mission time by stopping the robot from performing a socially risky motion. Consequently, this prevents accumulating social costs and leads to a better performance with respect to social metrics (but not necessarily $M_1 - M_2$) as shown in the rest of the results section.

A. Test Case R-I: Multi-Robot Single Human Real Tests

In a human-free environment Robot1 initially takes the task and when the human starts walking towards Robot1 the task is delegated to Robot2 through active coordination. Upon behavior change, non-adaptive replanning strategies cannot reassign the task to Robot1 due to the constraint of single attempt of active coordination, despite the replanning attempts. This happens while Robot1 is aware that it is the best candidate for the task.

Figure 5 shows the performance of the four different scenarios tested in reality. ARBR has superior performance compared to its non-adaptive counterpart across all scenarios. Similar social performance can be seen for the scenarios adopting the adaptive replanning strategy. However, the distance

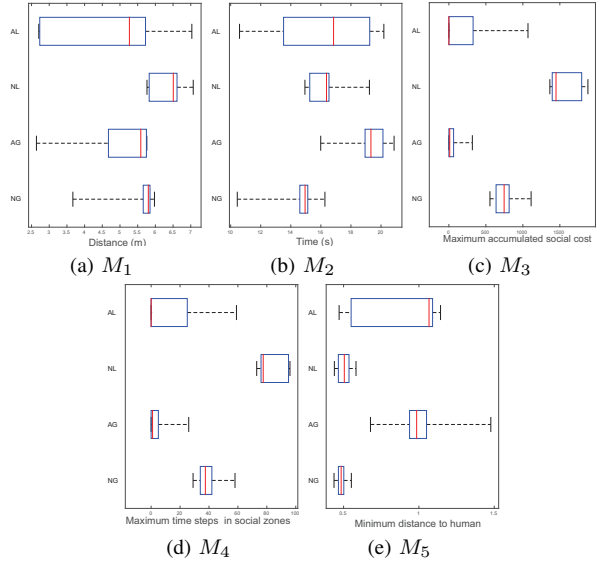


Fig. 5: Performance metrics for test case R-I obtained from 10 runs. The central red marks indicate the median, and the bottom and top edges of the box indicate the 25th and 75th percentiles, respectively.

and time ($M_1 - M_2$) for scenario *AG* are longer compared to scenario *AL*. This is because global perception is faster to detect the human and as a result, Robot2 is dispatched to the task at an earlier time and travels a larger distance before it is notified of the change and stops. As Robot1 has also stopped earlier, the remaining distance between Robot1 and the task is larger and requires more time to traverse. With local perception however, Robot1 is stopped later and as a result, has traveled a larger segment of its path to the task before resuming its motion for the second time.

The task plan executed in each scenario can be seen in Figure 6. We can observe that a different task assignment is done for scenarios with adaptive replanning compared to their non-adaptive counterparts. The waiting periods depicted in the task plots (yellow blocks) for Robot2 occur since despite the increasing risk of the task that is caused by the human behavior change, Robot2 is still the best candidate to take it. Nonetheless, the increasing risk forces the robot to stop and thus, this waiting behavior emerges. Finally, the risk gets large enough to notify Robot2 that Robot1 is the best candidate for taking the task.

The risk variation for scenarios *AG* and *AL* is depicted in Figure 7. We can see how the risk for the task initially increases for Robot1 in both plots. Consequently, a rebid for Robot1 happens and the task is delegated to Robot2. Robot2 predicts no risk for the task as long as the human is moving towards Robot1. Once the human changes his walking direction, the risk rises and rebidding occurs. As mentioned previously, multiple rebids prevent Robot2 from moving to the task. Finally, the risk is large enough to delegate the task to Robot1 by means of active coordination.

B. Test Case R-II: Multi-Robot Multi-Human Real Tests

In this test case, given the problem configuration in an empty arena, robot plans are as follows: Robot1 first takes “Task1” and then “Task3”, and Robot2 takes “Task2”. The challenging part of this test case is that Robot1 cannot observe the two humans that are about to start walking in the arena.

¹Instituto Portugues de Oncologia de Lisboa (IPOL)

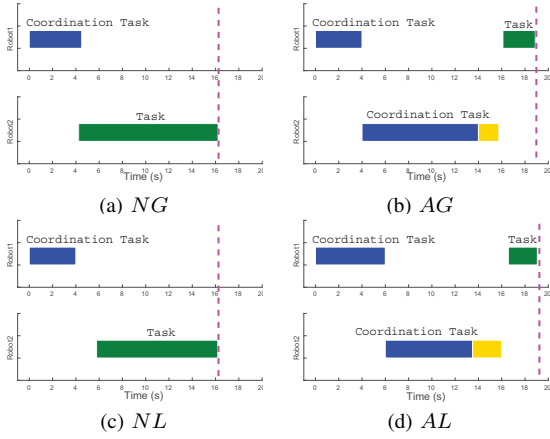


Fig. 6: Task assignment per robot over time for a sample run of test case R-I for scenarios *NG*, *AG*, *NL*, and *AL* respectively. End of mission (M_2) is marked by the vertical dashed line. Yellow blocks indicate the waiting periods.

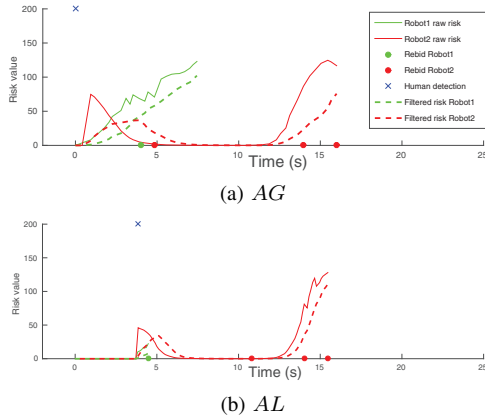


Fig. 7: Risk plots over time for Robot1 and Robot2 in test case R-I for a) scenario *AG*, and b) scenario *AL*.

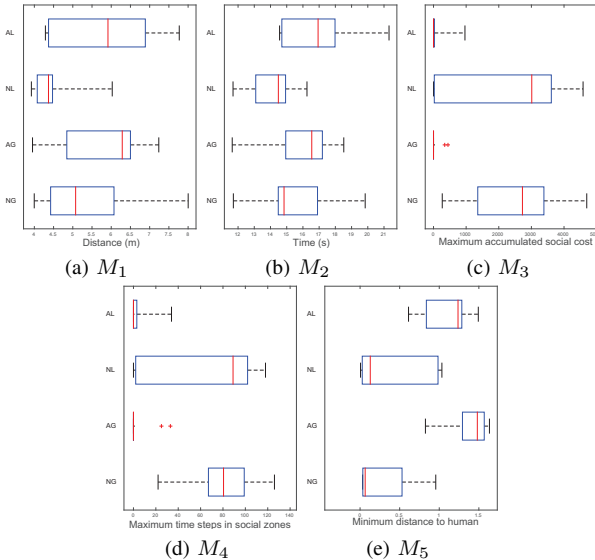


Fig. 8: Performance metrics for test case R-II obtained from 10 runs.

Given the different delays introduced in the starting time of each human’s walking motion, robots are faced with different social costs. Additionally, different perceptual information is available for decision making with local perception throughout runs depending on how the humans are relatively positioned with respect to the robots.

Fig. 8 shows the performance of the four different scenarios

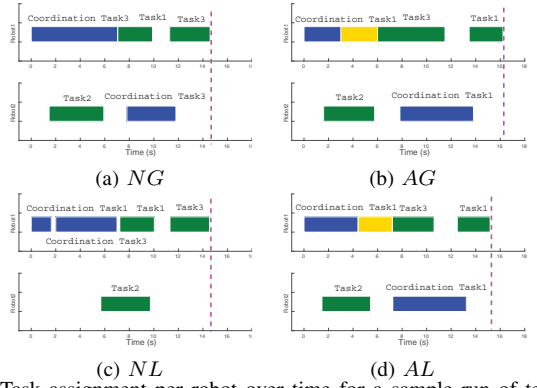


Fig. 9: Task assignment per robot over time for a sample run of test case R-II for scenarios *NG*, *AG*, *NL*, and *AL* respectively. End of mission (M_2) is marked by the vertical dashed line.

tested. Similar to the previous test cases, adaptive replanning has led to better performances with respect to social metrics ($M_3 - M_5$). For ($M_1 - M_2$) however, scenario *NG* and *NL* are performing better. We can also observe that global perception has led to slightly better performance in scenarios with adaptive replanning and contrarily, slightly worse performances in scenarios with non-adaptive replanning.

Fig. 9 shows the robot plans for a sample run of each scenario. We can see how the plans have changed when adopting ARBR. Moreover, we can see how waiting periods have been introduced to deal with increasing social costs. In this test case, in scenarios *AG* and *AL*, Robot1 initially moves to “Task1” it then stops when sensing the increased risk for all tasks. Later on it moves to “Task3” and then to “Task1”. In some runs where at least one human started to move after a larger delay, we observed either Robot2 initially taking “Task1” and then delegating it through active coordination, or completing “Task1”. This explains the larger variation in $M_1 - M_2$ for scenario *AG* and *AL* compared to scenarios *NG* and *NL*.

C. IPOL Simulations

Table II summarizes the results of our simulation suite. The columns of the table represent the number of robots ($\#R$), humans ($\#H$), tasks ($\#T$), the scenario and evaluation metrics, respectively. Human motion dynamics have been added to scenario labels. W stands for walking, S stands for static, and A stands for an alternating back and forth motion.

Looking at Table II we can see how the adaptive replanning method has led to a better performance with respect to social metrics $M_3 - M_5$, in all cases and furthermore, has succeeded in avoiding social costs in most cases. Additionally, the adaptive replanning method shows slightly better performance for $M_1 - M_2$. In general, this is not always the case, since avoiding social risks can sometimes lead to plans that incur larger costs in terms of distance and time.

To explain why ARBR performs better compared to its non-adaptive counterpart, consider the example where a human walks the width of a corridor with an alternating back and forth motion. As the human is constantly changing his walking behavior, the initial human motion prediction made by a non-adaptive approach can be incorrect. Thus, it will lead to a decision that is likely to be poor by the time the robot gets to a close vicinity of the human.

Results show that for the adaptive method, global perception has a better performance compared to local perception in the

experiments due to having a more accurate prediction of the human trajectory based on updated perceptual information. Robots with limited perception can only react to the human once the human is observed within the FOV of one of the robots. This typically, leads to later plan modification and occasionally, minor violations of social constraints.

On the contrary, we can see that for the non-adaptive method, local perception can performs better compared to global perception in terms of social metrics (see Table II scenarios A-NG and A-NL). The reason is that the replanning strategy of scenario A-NG relies on incorrect predictions in most cases, since the human behavior changes within the execution period of the plan and is no longer what was expected to be at the decision making time. Furthermore, due to having global perception, A-NG is comparably more reactive to change.

VII. CONCLUSION AND FUTURE WORK

Adaptive risk-based replanning has shown to have superior performance in terms of social metrics in all test cases studied in this paper. This comes with the price of longer plans in terms of traveled distance and time in some cases. Global perception has shown to improve the performance of the robot team that has an adaptive replanning strategy. However, without adaptation to social risks and changes in the environment, global perception can lead to worse performances compared to local perception. ARBR with local perception shows comparable performance to global perception. As the local perception gets more restricted in a highly dynamic and stochastic environment, the performance of ARBR strategy degrades. However, a realistic conic FOV of 4 m range and 65° is shown to do very well in our experiments with stochastic human behavior. As the environment gets more dynamic and stochastic, ARBR still performs significantly better compared to its non-adaptive

counterpart but we can observe that social constraints are violated from time to time.

For further improvements the following points can be considered. The level of risk that a robot will take before reconsidering its decision can give a conservative or risk-taking character to the robot. This is a choice that should be made based on the specific environment that the robots are operating in, either by the designer or learned from available data for different situations. Additionally, the risk formulation can be improved by including spatio-temporal and map-based information in the next steps.

REFERENCES

- [1] T. Kruse, A. K. Pandey, R. Alami, and A. Kirsch, "Human-aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 61, no. 12, pp. 1726 – 1743, 2013.
- [2] K. Charalampous, I. Kostavelis, and A. Gasteratos, "Recent trends in social aware robot navigation: A survey," *Robotics and Autonomous Systems*, vol. 93, pp. 85–104, 2017.
- [3] G. A. Korsah, A. Stentz, and M. B. Dias, "A comprehensive taxonomy for multi-robot task allocation," *The International Journal of Robotics Research*, vol. 32, no. 12, pp. 1495–1512, 2013.
- [4] M. B. Dias, R. Zlot, N. Kalra, and A. Stentz, "Market-based multi-robot coordination: A survey and analysis," *Proceedings of the IEEE*, vol. 94, no. 7, pp. 1257–1270, 2006.
- [5] Z. Talebpour and A. Martinoli, "Risk-based human-aware multi-robot coordination in dynamic environments shared with humans," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 3365–3372, 2018.
- [6] E. Nunes, M. Manner, H. Mitiche, and M. Gini, "A taxonomy for task allocation problems with temporal and ordering constraints," *Robotics and Autonomous Systems*, vol. 90, pp. 55–70, 2017.
- [7] Z. Talebpour and A. Martinoli, "Multi-robot coordination in dynamic environments shared with humans," in *IEEE International Conference on Robotics and Automation*, pp. 1105–1112, 2018.
- [8] Y. Cao, W. Yu, W. Ren, and G. Chen, "An overview of recent progress in the study of distributed multi-agent coordination," *IEEE Transactions on Industrial Informatics*, vol. 9, no. 1, pp. 427–438, 2013.
- [9] J. Rios-Martinez, A. Spalanzani, and C. Laugier, "Understanding human interaction for probabilistic autonomous navigation using risk-RRT approach," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 2014–2019, 2011.
- [10] E. T. Hall, "The hidden dimension," 1966.
- [11] Z. Talebpour, D. Viswanathan, R. Ventura, G. Englebienne, and A. Martinoli, "Incorporating perception uncertainty in human-aware navigation: A comparative study," in *IEEE Int. Symp. on Robot and Human Interactive Communication*, pp. 570–577, 2016.
- [12] R. Ventura and A. Ahmad, "Towards optimal robot navigation in domestic spaces," in *RoboCup 2014: Robot World Cup XVIII*, pp. 318–331, Springer, 2015.
- [13] J. V. Gómez, N. Mavridis, and S. Garrido, "Fast marching solution for the social path planning problem," in *IEEE International Conference on Robotics and Automation*, pp. 2243–2248, 2014.
- [14] S. Sarel-Talay, T. R. Balch, and N. Erdogan, "Multiple traveling robot problem: A solution based on dynamic task selection and robust execution," *IEEE/ASME Transactions on Mechatronics*, vol. 14, pp. 198–206, April 2009.
- [15] B. P. Sellner and R. Simmons, "Towards proactive replanning for multi-robot teams," in *Proceedings of the 5th International Workshop on Planning and Scheduling in Space 2006*, October 2006.
- [16] N. Buckman, H.-L. Choi, and J. P. How, "Partial replanning for decentralized dynamic task allocation," in *AIAA Scitech Forum*, 2019.
- [17] N. Kalra, D. Ferguson, and A. Stentz, "Hoplites: A market-based framework for planned tight coordination in multi-robot teams," in *IEEE International Conference on Robotics and Automation*, pp. 1170–1177, 2005.
- [18] Z. Talebpour, S. Savarè, and A. Martinoli, "Market-based coordination in dynamic environments based on the Hoplites framework," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pp. 1105–1112, 2017.
- [19] J. Sequeira, P. Lima, A. Saffiotti, V. Gonzalez-Pacheco, and M. Salichs, "MONarCH: Multi-robot cognitive systems operating in hospitals," in *ICRA workshop on many robot systems*, 2013.
- [20] J. Messias, R. Ventura, P. Lima, J. Sequeira, P. Alvito, C. Marques, and P. Carriço, "A robotic platform for edutainment activities in a pediatric hospital," in *IEEE International Conference on Autonomous Robot Systems and Competitions*, pp. 193–198, 2014.

TABLE II: Summarized experimental results for IPOL simulations. The median and mean are reported for each metric respectively. See text for legend.

#R	#H	#T	Scenario	M_1	M_2	M_3	M_4	M_5
2	2	1	W-NL	15.42	36.01	0	0	1.12
				15.93	36.89	105	3.6	1.02
2	2	1	W-AL	14.93	35	0	0	1.5
				14.32	35.11	0	0	1.73
2	1	1	S-NG	18	46.07	0	0	2.63
				19.51	47.52	433	28.2	2.14
2	1	1	S-NL	19.82	50.05	354	10	0.64
				19.83	51.47	1865	68.14	0.71
2	1	1	S-AG	18.21	44.08	0	0	2.26
				18.19	43.89	0	0	2.31
2	1	1	S-AL	18.48	50.07	0	0	2.27
				18.45	49.14	0	0	2.29
2	12	1	W-NG	11.80	26.24	5800	220	0.51
				13.11	30.84	7560	318	0.73
2	12	1	W-NL	14.22	38.57	18000	605	0.46
				14.35	38.64	17212	598	0.38
2	12	1	W-AG	10.22	26.31	0	0	1.81
				10.41	35.88	881	54	1.55
2	12	1	W-AL	10.11	31.18	4941	201	0.68
				13.02	45.54	4774	188	1.07
3	1	5	A-NG	48.12	79.98	48	3	0.84
				52.44	99.71	112	11.47	0.97
3	1	5	A-NL	44.87	75.11	0	0	1.87
				47.55	81.73	72	12.66	1.17
3	1	5	A-AG	37.05	61.55	0	0	1.95
				36.88	64.61	5	0.3	1.83
3	1	5	A-AL	40.21	77.37	0	0	2.01
				40.24	71.48	31	7.1	1.92