Distributed Plan Reconfiguration via Knowledge Transfer in Multi-agent Systems under Local LTL Specifications

Meng Guo and Dimos V. Dimarogonas

Abstract—We propose a cooperative motion and task planning scheme for multi-agent systems where the agents have independently-assigned local tasks, specified as Linear Temporal Logic (LTL) formulas. These tasks contain hard and soft sub-specifications. A least-violating initial plan is synthesized first for the potentially infeasible task and the partially-known workspace. While the system runs, each agent updates its knowledge about the workspace via its sensing capability and shares this knowledge with its neighboring agents. Based on this update, each agent verifies and revises its plan in real time. It is ensured that the hard specification is always fulfilled and the satisfaction for the soft specification is improved gradually. The design is distributed as only local interactions are assumed. The overall framework is demonstrated by a case study.

I. INTRODUCTION

Temporal-logic-based motion planning has gained significant attention in recent years, as it provides a correct-bydesign controller synthesis approach for autonomous robots. Temporal logics such as Linear Temporal Logic (LTL) and Computation Tree Logic (CTL) provide formal high level languages that can describe planning objectives more complex than the well-studied point-to-point navigation [20]. The task is given as a logic formula with respect to the discretized abstraction of the robot motion modelled as a finite transition system [1], [5]. Then a high-level discrete plan is found by off-the-shelf model-checking algorithms given the finite transition system and the task specification [3], [9]. This discrete plan is then implemented through the low-level hybrid controller [17], [19]. This methodology has also been applied for multi-agent systems [8], [15], [26]. Most of the existing work focuses on how to decompose a global specification to bisimilar local ones in a top-down approach, which can be then implemented by individual agents in a synchronized [16] and partially-synchronized [17] manner.

Here we instead assume that local task specifications are assigned independently to each agent and there is no prespecified global task [10], [13]. Since the workspace is only partially-known, the task specification might be infeasible given the initial workspace model and the motion plan might need to be revised whenever the workspace model is updated. Our previous work [13] addresses the first aspect by synthesizing the motion plan that fulfills a potentially infeasible task the most. In this paper, we further improve this technique by considering hard and soft sub-specifications, where the hard part should be always fulfilled while the soft part can be relaxed. Similar problems are discussed in [23], [24] to find the least-violating strategy under a set of safety rules. However the level of satisfiability is measured in a different way from our approach. In particular, we do not only measure how many states along the plan violate the task, but also how much each of those states violates the task. Moreover we propose a weight function that balances the task satisfiability of a plan and its implementation cost. Regarding the second aspect, related approaches can be found, e.g., in our earlier work [12] by real-time plan revising, in [21] by local "patching" and in [22] by workspace re-abstraction. In this paper, a new technique is proposed taking into account the fact that the knowledge update comes not only from local sensing but also from interagent communications.

The proposed scheme has the following attributes: (1) locally assigned tasks and a partially-known workspace are considered; (2) it is distributed since only local interaction between neighboring agents is assumed; (3) the communication payload is significantly reduced compared to a fully synchronized solution; (4) while the hard specification that preserves safety is always guaranteed by the real-time reconfiguration algorithm, the satisfaction of the soft specification is improved by triggering the optimal synthesis algorithm in an event-based fashion; (5) it can potentially be applied to many existing partition and motion planning techniques, such as probabilistic roadmap method [8], [14], rapidly-exploring random trees [1], [2], and navigation functions [18].

The rest of the paper is organized as follows: Section II briefly introduces essential preliminaries. In Section III, we describe how to synthesize an initial optimal plan. Section IV focuses on how the knowledge is transferred among the agents. Real-time algorithms to update the motion plan are provided in Section V. Section VI summarizes the overall structure and the case study is presented in Section VII.

II. PRELIMINARY

Given a team of heterogeneous agents $k \in \mathcal{K} = \{1, 2\cdots, K\}$, with unique identities (IDs) k, they move within a workspace Π which consists of N regions with unique and pre-defined labels π_n , $n = 1, 2\cdots, N$.

A. Local Task Specifications

The basic ingredients of an Linear Temporal Logic (LTL) formula φ are a set of atomic propositions (APs) *AP* and several boolean and temporal operators. We omit the syntax and semantics of LTL here and refer the readers to

The authors are with the KTH Centre for Autonomous Systems and ACCESS Linnaeus Center, EES, KTH Royal Institute of Technology, SE-100 44, Stockholm, Sweden. mengg, dimos@kth.se. This work was supported by the Swedish Research Council (VR) and EU STREP RECONFIG: FP7-ICT-2011-9-600825.

Chapter 5 of [3]. There is a union of infinite words that satisfy φ : Words $(\varphi) = \{\sigma \in (2^{AP})^{\omega} | \sigma \models \varphi\}$, where $\models \subseteq (2^{AP})^{\omega} \times \varphi$ is the satisfaction relation [7].

Furthermore there exists a Nondeterministic Büchi automaton (NBA) \mathcal{A}_{φ} over 2^{AP} corresponding to φ [3], [7]. It is defined as $\mathcal{A}_{\varphi} = (Q, 2^{AP}, \delta, Q_0, \mathcal{F})$ where Q is a finite set of states; $Q_0 \subseteq Q$ is the set of initial states, 2^{AP} is the set of input alphabets; $\delta : Q \times 2^{AP} \to 2^Q$ is a transition relation and $\mathcal{F} \subseteq Q$ is a set of accepting states. Denote by $\chi(q_m, q_n) = \{l \in 2^{AP} | q_n \in \delta(q_m, l)\}$ the set of all input alphabets that enable the transition from q_m to q_n . An accepting run of \mathcal{A}_{φ} is an infinite run that intersects with \mathcal{F} infinitely often. Denote by $\mathcal{L}_{\omega}(\mathcal{A}_{\varphi})$ the accepted language of \mathcal{A}_{φ} , It holds that $\operatorname{Words}(\varphi) = \mathcal{L}_{\omega}(\mathcal{A}_{\varphi})$ [3]. There are fast translation algorithms [11] from an LTL formula to NBA. This process can be done in time and space $2^{\mathcal{O}(|\varphi|)}$.

Denote by $AP_k = \{a_1, a_2, \dots, a_{|AP_k|}\}$ the set of APs known to agent k. Its task specification φ_k is an LTL formula over AP_k . We assume that φ_k remains unchanged for all agents once the system starts. Denote by $\varphi_k|_{AP_k}$ as the set of APs appearing in φ_k . In particular, we consider the task specification φ_k with the following structure:

$$\varphi_k = \varphi_k^{\text{soft}} \land \varphi_k^{\text{hard}},\tag{1}$$

where φ_k^{soft} and φ_k^{hard} are "soft" and "hard" sub-formulas over AP_k . φ_k^{hard} could include safety constraints like collisionavoidance: "avoid all obstacles" or energy-supply guarantee: "visit the charging station infinitely often". Introducing soft and hard specifications is due to the observation that the partially-known workspace might render parts of the specification infeasible initially and thus yielding the needs for them to be relaxed, while the safety-critical parts should not be relaxed during the process.

B. Agent Description and Partially-known Workspace

The motion of agent k in the workspace is described by a deterministic finite-transition system (FTS):

$$\mathcal{T}_k = (\Pi, \longrightarrow_k, \Pi_{0,k}, AP_k, L_k, W_k), \tag{2}$$

where (i) $\longrightarrow_k \subseteq \Pi \times \Pi$ is the transition relation; (ii) $\Pi_{0,k} \subseteq \Pi$ is the set of initial regions where agent k may start; (iii) $L_k : \Pi \rightarrow 2^{AP_k}$ is the labelling function that represents agent k's knowledge. $L_k(\pi_i)$ is the set of APs satisfied by π_i ; (iv) $W_k : \longrightarrow_k \rightarrow \mathbb{R}^+$ is the cost of each transition.

We assume that each agent has only partial knowledge of the workspace initially. Thus \mathcal{T}_k might be updated afterwards as described in Section IV-B. Denote by \mathcal{T}_k^t the FTS of agent k at time t. Note that AP_k might be different among the agents due to heterogeneity; \mathcal{T}_k does not have a terminal state [3]. An infinite path of \mathcal{T}_k is an infinite sequence of states $\tau_k = \pi_0 \pi_1 \pi_2 \cdots$ such that $(\pi_i, \pi_{i+1}) \in \longrightarrow_k$ for all i > 0. Its trace is defined as the sequence of APs that are true at the states along the path, i.e., $\text{trace}(\tau_k) =$ $L_k(\pi_0)L_k(\pi_1)L_k(\pi_2)\cdots$. Since φ_k^{hard} and φ_k^{soft} are LTL formulas over AP_k , then $\text{trace}(\tau_k) \models \varphi_k^{\text{hard}}$ if and only if $\text{trace}(\tau_k) \in \text{Words}(\varphi_k^{\text{hard}})$ and $\text{trace}(\tau_k) \models \varphi_k^{\text{soft}}$ if and only if $\text{trace}(\tau_k) \in \text{Words}(\varphi_k^{\text{soft}})$. Definition 1: Given an infinite path $\tau_k = \pi_0 \pi_1 \pi_2 \dots$ of \mathcal{T}_k and its task φ_k, τ_k is called: (i) valid if $(\pi_i, \pi_{i+1}) \in \longrightarrow_k$, for $i = 0, 1, 2 \dots$; (ii) safe if $\text{trace}(\tau_k) \models \varphi_k^{\text{hard}}$; (iii) satisfying if $\text{trace}(\tau_k) \models \varphi_k$.

III. INITIAL OPTIMAL PLAN SYNTHESIS

Denote by $\mathcal{A}_k^{\text{hard}} = (Q_1, 2^{AP_k}, \delta_1, Q_{1,0}, \mathcal{F}_1)$ and $\mathcal{A}_k^{\text{soft}} = (Q_2, 2^{AP_k}, \delta_2, Q_{2,0}, \mathcal{F}_2)$ as the NBA associated with φ_k^{hard} and φ_k^{soft} , respectively. The functions $\chi_1()$ of $\mathcal{A}_k^{\text{hard}}$ and $\chi_2()$ of $\mathcal{A}_k^{\text{soft}}$ are defined analogously as in Section II-A. Now we propose a way to construct the intersection of $\mathcal{A}_k^{\text{hard}}$ and $\mathcal{A}_k^{\text{soft}}$, which is safety-ensured and relaxed.

Definition 2 (Relaxed Automata Intersection): The relaxed intersection of \mathcal{A}_k^{hard} and \mathcal{A}_k^{soft} is defined by:

$$\tilde{\mathcal{A}}_{\varphi_k} = (Q, 2^{AP_k}, \delta, Q_0, \mathcal{F}), \tag{3}$$

where $Q = Q_1 \times Q_2 \times \{1,2\}$; $Q_0 = Q_{1,0} \times Q_{2,0} \times \{1\}$; $\mathcal{F} = \mathcal{F}_1 \times Q_2 \times \{1\}$; $\delta : Q \times 2^{AP_k} \to 2^Q$, with $\langle \check{q}_1, \check{q}_2, \check{t} \rangle \in \delta(\langle q_1, q_2, t \rangle, l)$ when the following three conditions hold: (1) $l \in \chi_1(q_1, \check{q}_1)$; (2) $\chi_2(q_2, \check{q}_2) \neq \emptyset$; (3) $q_t \notin \mathcal{F}_t$ and $\check{t} = t$, or $q_t \in \mathcal{F}_t$ and $\check{t} = \mod(t, 2) + 1$, where $t \in \{1, 2\}$ and mod is the modulo operation.

Note we relax the requirement that there should exist a common input alphabet that enables the transitions from q_i to \check{q}_i for both i = 1, 2, compared with the standard definition of Büchi automata intersection (see Chapter 4.3 of [3]). An accepting run R of $\tilde{\mathcal{A}}_{\varphi_k}$ intersects with \mathcal{F} infinitely often. The last component $t \in \{1, 2\}$ in Q ensures that R has to intersect with both $\mathcal{F}_1 \times Q_2 \times \{1\}$ and $Q_1 \times \mathcal{F}_2 \times \{2\}$. Denote by $R|_{Q_1}$ the projection of R onto the states of \mathcal{A}_k^{hard} .

Theorem 1: Given an accepting run R of \mathcal{A}_{φ_k} , $R|_{Q_1}$ is an accepting run of $\mathcal{A}_k^{\text{hard}}$. Moreover, $\mathcal{L}_{\omega}(\tilde{\mathcal{A}}_{\varphi_k}) \subseteq \mathcal{L}_{\omega}(\mathcal{A}_k^{\text{hard}})$.

Proof: The projection of \mathcal{F} onto Q_1 is \mathcal{F}_1 , therefore one of the accepting states in \mathcal{F}_1 is visited infinitely often by $R|_{Q_1}$. Since $l \in \chi_1(q_1, \check{q}_1)$ is ensured by the definition of δ , $R|_{Q_1}$ is an accepting run of \mathcal{A}_k^{hard} . Secondly, given any word $\sigma \in \mathcal{L}_{\omega}(\tilde{\mathcal{A}}_{\varphi_k})$, σ results in an accepting run of $\tilde{\mathcal{A}}_{\varphi_k}$, denoted by R_{σ} . Since $R_{\sigma}|_{Q_1}$ is an accepting run of \mathcal{A}_k^{hard} , it implies that $\sigma \in \mathcal{L}_{\omega}(\mathcal{A}_k^{hard})$. Thus for any $\sigma \in \mathcal{L}_{\omega}(\tilde{\mathcal{A}}_{\varphi_k})$, $\sigma \in \mathcal{L}_{\omega}(\mathcal{A}_k^{hard})$ holds, namely $\mathcal{L}_{\omega}(\tilde{\mathcal{A}}_{\varphi_k}) \subseteq \mathcal{L}_{\omega}(\mathcal{A}_k^{hard})$. \blacksquare Since we need to guarantee that φ_k^{hard} is fulfilled fully and

Since we need to guarantee that φ_k^{soft} is fulfilled fully and φ_k^{soft} is fulfilled the most, the standard model-checking-based algorithm [8], [9] may fail to provide any solution. Thus we rely product automaton construction proposed in [13].

Definition 3 (Weighted Product Automaton): The

weighted product Büchi automaton $\tilde{\mathcal{A}}_{p,k} = \mathcal{T}_k \times \tilde{\mathcal{A}}_{\varphi_k} = (Q', \delta', Q'_0, \mathcal{F}', W_p)$ is defined as follows:

- $Q' = \Pi \times Q$ and $q' = \langle \pi, q \rangle$, $\forall \pi \in \Pi$ and $\forall q \in Q$.
- $\delta' : Q' \to 2^{Q'}$. $\langle \pi_j, q_n \rangle \in \delta'(\langle \pi_i, q_m \rangle)$ iff $(\pi_i, \pi_j) \in \longrightarrow_k$ and $q_n \in \delta(q_m, L_k(\pi_i))$.
- $Q'_0 = \Pi_{0,k} \times Q_0$ is the set of initial states.
- $\mathcal{F}' = \Pi \times \mathcal{F}$ is the set of accepting states.
- W_p : $Q' \times Q' \to \mathbb{R}^+$ is the weight function, defined as $W_p(\langle \pi_i, q_m \rangle, \langle \pi_j, q_n \rangle) = W_k(\pi_i, \pi_j) + \alpha \cdot \text{Dist}(L_k(\pi_i), \chi_2(q_2, \check{q}_2))$, where $\langle \pi_j, q_n \rangle \in \delta'(\langle \pi_i, q_m \rangle)$, $q_m = \langle q_1, q_2, t \rangle$ and $q_n = \langle \check{q}_1, \check{q}_2, \check{t} \rangle$; $\alpha \geq 0$ is a design parameter.

We denote by $\mathcal{A}_{p,k}^t$ as the product automaton corresponding to \mathcal{T}_k^t at time t. The detailed definition and examples of computing Dist() can be found in Section III.A of [13]. We briefly restate it here. Firstly, an evaluation function over 2^{AP_k} is defined: Eval: $2^{AP_k} \rightarrow \{0, 1\}^{|AP_k|}$. Eval(l) = $\nu \iff [\nu_i] = 1$, if $a_i \in l$ and $[\nu_i] = 0$, if $a_i \notin l$, where $l \in 2^{AP_k}, \nu \in \{0, 1\}^{|AP_k|}$ and $i = 1, 2 \cdots, |AP_k|$. Then a metric $(2^{AP_k}, \rho)$ is defined as $\rho(l, l') = \|\nu - \nu'\|_1 =$ $\sum_{i=1}^{|AP_k|} |\nu_i - \nu'_i|$, where $\nu = \text{Eval}(l), \nu' = \text{Eval}(l')$ and $l, l' \in 2^{AP_k}$. $\|\cdot\|_1$ is the ℓ_1 norm [6]. The distance between an element $l \in 2^{AP_k}$ to a set $\chi \subseteq 2^{AP_k} (\chi \neq \emptyset)$ is given by [6]: $\text{Dist}(l, \chi) = 0$, if $l \in \chi$ and $\text{Dist}(l, \chi) =$ $\min_{l' \in \chi} \rho(l, l')$ otherwise. Remind that $\chi_2(q_2, \check{q}_2) \neq \emptyset$ by the definition of δ in Definition 2. Moreover, it is easy to see that $\text{Dist}(L_k(\pi_i), \chi_1(q_1, \check{q}_1)) = 0, \forall \langle \pi_i, q_n \rangle \in$ $\delta'(\langle \pi_i, q_m \rangle)$ from the definition of δ' .

The weight function W_p in δ' consists of two parts: $W_k(\pi_i, \pi_j)$ measures the implementation cost of the transition from π_i to π_i ; Dist $(L_k(\pi_i), \chi_2(q_2, \check{q}_2))$ measures how much this transition violates the constraints imposed by $\mathcal{A}_k^{\text{soft}}$. The design parameter α reflects the relative penalty on violating the soft specification, which should be chosen relatively large under partially-known workspace.

Theorem 2: Assume R_k is an accepting run of $\mathcal{A}_{p,k}$. Then $\tau_k = R_k|_{\Pi}$, is both valid and safe for \mathcal{T}_k and φ_k .

Proof: The fact that τ_k is valid can be verified from the definition of δ' . Because every transition in δ' when projected onto Π is a valid transition within \longrightarrow_k , meaning that τ_k is always implementable by \mathcal{T}_k . Secondly, since R_k is an accepting run of $\mathcal{A}_{p,k}$, then $trace(\tau_k) \in \mathcal{L}_{\omega}(\mathcal{A}_{\varphi_k})$, which implies $trace(\tau_k) \in \mathcal{L}_{\omega}(\mathcal{A}_k^{hard})$ by Theorem 1. Since $\operatorname{Words}(\varphi_k^{\operatorname{hard}}) = \mathcal{L}_{\omega}(\mathcal{A}_k^{\operatorname{hard}}), \operatorname{trace}(\tau_k) \in \operatorname{Words}(\varphi_k^{\operatorname{hard}}),$ which indicates that τ_k is also safe by Definition 1.

In order to measure the implementation cost of different accepting runs of $A_{p,k}$ and how much they violate the soft specification, we consider the accepting runs with the following *prefix-suffix* structure: R_k = $q'_0 q'_1 \cdots q'_{f-1} [q'_f q'_{f+1} \cdots q'_n]^{\omega}$, where $q'_0 \in Q'_0$ and $q'_f \in \mathcal{F}'$. The prefix part $(q'_0 q'_1 \cdots q'_{f-1} q'_f)$ of R_k from an initial state q'_0 to one accepting state q'_f is executed only once. The suffix part $(q'_f q'_{f+1} \cdots q'_n)$ of R_k from q'_f back to itself is repeated infinitely. The total cost of R_k is defined by: $\operatorname{Cost}(R_k, \tilde{\mathcal{A}}_{p,k}) = \sum_{i=0}^{f-1} W_p(q_i', q_{i+1}') +$ $\gamma \sum_{i=1}^{n-1} W_p(q'_i, q'_{i+1})$, where the design parameter $\gamma \ge 0$ represents the relative weighting on the cost of transient response (the prefix) and steady response (the suffix) to the task specification [26]. The accepting run R_k of $\mathcal{A}_{p,k}$ that minimizes $Cost(R_k, \mathcal{A}_{p,k})$ is called the optimal accepting run and denoted by $R_{\text{opt},k}$. Its projection onto Π , $\tau_{\text{opt},k} =$ $R_{\text{opt},k}|_{\Pi}$, is called the optimal motion plan.

Given the desired value of α and γ , Algorithm 1 in [13] generates the optimal accepting run of $\mathcal{A}_{p,k}$ and the corresponding plan in the prefix-suffix format. We omit the complete algorithm here due to limited space. Briefly speaking, it utilizes Dijkstra's algorithm [20] for computing the shortest path from any initial state to an accepting state and the the shortest cycle containing this accepting state. This procedure

Algorithm 1: Optimal Plan Synthesis OptiPlan()

Input: $\mathcal{T}_k, \, \mathcal{A}_{\varphi_k}$ **Output**: $R_{\text{opt},k}$, $\tau_{\text{opt},k}$ 1 Choose the desired α and γ ; 2 Compute $\tilde{\mathcal{A}}_{p,k} = \mathcal{T}_k \times \tilde{\mathcal{A}}_{\varphi_k}$ by Definition 3;

- 3 Derive $R_{\text{opt},k}$ in $\overline{A}_{p,k}$ by Algorithm 1 in [13];
- 4 $\tau_{\text{opt},k} = R_{\text{opt},k}|_{\Pi}$

is summarized in Algorithm 1. By Theorem 2, $\tau_{opt,k}$ is always valid and safe no matter how α and γ are chosen.

IV. KNOWLEDGE UPDATE AND TRANSFER

The agents have both the sensing ability to discover the workspace and the communication functionality to share knowledge with their neighboring agents.

Denote by Sense^t_k = { $A(t), A_{\neg}(t), B(t)$ } as the set of sensing information [12] obtained at time $t \ge 0$ by agent k, where: Type-A information: $((\pi_i, \pi_j), w) \in A(t)$ if a new transition (π_i, π_j) is added to \longrightarrow_k^t with weight w, or the weight of an existing transition (π_i, π_i) is updated to w. $(\pi_i, \pi_i) \in A_{\neg}(t)$ if (π_i, π_i) needs to be removed from \longrightarrow_k^t . Type-B information: denote $(a, \Pi_a, \Pi_{a,\neg}) \in B(t)$ where $a \in AP_k$; $\Pi_a \subseteq \Pi$ is the set of regions that satisfy a; $\Pi_{a,\neg} \subseteq \Pi$ is the set of regions that do not satisfy a. This sensing functionality can be modelled by assigning a sensing radius $S_k \ge 0$, such that all regions lying in the sphere $\{\mathbf{p} \in \mathbb{R}^n \mid |\mathbf{p} - \mathbf{p}_k^t| \le S_k\}$ are visible, where $\mathbf{p}_k^t \in \mathbb{R}^n$ is agent k's position at time t.

A. Knowledge Transfer

In this section, we explain how the knowledge about the workspace is shared among the agents.

1) Communication Network: It represents how the information flows among the agents. Each agent k has a set of neighboring agents, denoted by $\mathcal{K}_k \subseteq \mathcal{K}$. Agent k can send messages directly to any agent belonging to \mathcal{K}_k . We take into account two different ways to model the communication network: (1) global communication with a fixed topology; (2) limited communication with a dynamic topology.

2) Communication Protocol: Agent k is interested in all the propositions appearing in φ_k , namely $\varphi_k|_{AP_k}$. We propose a subscriber-publisher mechanism to reduce the communication load. Whenever agent g communicates with agent $k \in \mathcal{K}_a^t$ for the *first*, it follows the *subscribing* procedure: agent g sends a request message to agent k that **Request** $_{q,k}^{t} = \varphi_{g}|_{AP_{q}}$, which informs agent k the set of propositions agent g is interested. Each agent has a subscriber list, containing the requests it has received. Note that each agent sends a request to any of its neighbors only once since it also keeps track of the agents which it has subscribed to.

Afterwards, the publishing phase of each agent follows an event-driven approach: whenever an agent k has obtained a sensing update \mathbf{Sense}_{k}^{t} , it checks its subscriber list whether the content might be of interest to any of the subscribers regarding some propositions. If it is of interest to agent q,

Algorithm	2:	Transfer	Knowledge,	TranKnow())
-----------	----	----------	------------	------------	---

Input: Sense^t_k, Request^{t₀}_{g,k} Output: Reply^t_{k,g} 1 forall the Request^{t₀}_{g,k} received at t₀ < t do 2 if $g \in \mathcal{K}_k^t$ then 3 forall the $a \in \varphi_g|_{AP_g}$ do 4 if $(a, \tilde{\Pi}_a, \tilde{\Pi}_{a,\neg}) \in B(t)$ of Sense^t_k then 5 if $(a, \tilde{\Pi}_a, \tilde{\Pi}_a, \tilde{\Pi}_{a,\neg})$ to Reply^t_{k,g}

for example regarding proposition a, then agent k checks if $g \in \mathcal{K}_k^t$. If so, it publishes a reply message to agent g that $\mathbf{Reply}_{k,g}^t = \{(a, \Pi_a, \Pi_{a,\neg})\}$, where $a \in \varphi_g|_{AP_g}$; $\Pi_a, \Pi_{a,\neg} \subseteq \Pi$ are defined as in B(t). The above procedure is summarized in Algorithm 2. Note that through this communication mechanism, any request only needs to be sent once and every reply message contains useful knowledge.

B. Knowledge Update

At time t, agent k might obtain new knowledge from **Sense**^t_k and **Reply**^t_{g,k}, based on which it needs to update its FTS accordingly. Denote by $\mathcal{T}_{k}^{t^{-}}$ and $\mathcal{T}_{k}^{t^{+}}$ as the FTS before and after the update at time t. Given **Sense**^t_k, $\mathcal{T}_{k}^{t^{+}}$ is obtained by updating the transition relation of $\mathcal{T}_{k}^{t^{-}}$ according to $A(t), A_{\neg}(t)$ and the labelling function of $\mathcal{T}_{k}^{t^{-}}$ according to B(t). Given $(a, \tilde{\Pi}_{a}, \tilde{\Pi}_{a,\neg}) \in \mathbf{Reply}_{g,k}^{t}$, for regions inside $\tilde{\Pi}_{a}$, if proposition a does not belong to $L_{k}^{t^{-}}(\pi)$, a is added to $L_{k}^{t^{-}}(\pi)$ and π is added to $\tilde{\Pi}_{k}^{t}$, where $\tilde{\Pi}_{k}^{t} \subseteq \Pi$ is used to store the regions of which the labelling function is changed. Similar procedure is applied to $\tilde{\Pi}_{a,\neg}$; $\tilde{\Pi}_{k}^{t}$ is used to update the product automaton later. Note that if **Sense**^t_k and **Reply**^t_{g,k} are empty, $\mathcal{T}_{k}^{t^{+}}$ remains the same as $\mathcal{T}_{k}^{t^{-}}$.

V. REAL-TIME RECONFIGURATION

Since \mathcal{T}_k^t might be updated as described in Section IV, the optimal accepting run and its corresponding optimal motion plan derived by Algorithm 1 in Section III need to be evaluated regarding its validity, safety and optimality.

A. Product Automaton Update

Denote by $\tilde{\mathcal{A}}_{p,k}^{t^-}$ and $\tilde{\mathcal{A}}_{p,k}^{t^+}$ as the product automaton before and after the update at time *t*. It is possible to reconstruct completely $\tilde{\mathcal{A}}_{p,k}^{t^+}$ by Definition 3 using $\tilde{\mathcal{A}}_{\varphi_k}$ and the updated $\mathcal{T}_k^{t^+}$. Here we propose to revise $\tilde{\mathcal{A}}_{p,k}^{t^-}$ locally based on the latest changes in \mathcal{T}_k^t .

Definition 4: Given $\mathcal{T}_k^{t^+}$, $(\langle \pi_i, q_m \rangle, \langle \pi_j, q_n \rangle) \in \delta_k^{\prime t^-}$ is called: (i) *invalid* if $(\pi_i, \pi_j) \notin \longrightarrow_k^{t^+}$; (ii) *unsafe* if $L_k^{t^+}(\pi_i) \notin \chi_1(q_m|_{Q_1}, q_n|_{Q_1})$.

To update $\tilde{\mathcal{A}}_{p,k}^{t^-}$, all transitions in $\tilde{\mathcal{A}}_{p,k}^{t^-}$ corresponding to removed transitions contained in $A_{\neg}(t)$ have to be removed. On the other hand, all transitions in $\tilde{\mathcal{A}}_{p,k}^{t^-}$ corresponding to added or modified transitions contained in A(t) and regions with modified labelling function contained in $\tilde{\Pi}_k^t$ need to be

Algorithm 3: Revise the plan, ReviPlan()

Input: $R_k^{t^-}, \Xi_k^t, \aleph_k^t, \tilde{\mathcal{A}}_{p,k}^{t^+}$ Output: $R_k^{t^+}, \tau_k^{t^+}$ 1 $R_k^{t^+} = \text{Revise}(\tilde{\mathcal{A}}_{p,k}^{t^+}, R_k^{t^-}, \Xi_k^t \cup \aleph_k^t)$ by Algorithm 5 in [12]; 2 $\tau_k^{t^+} = R_k^{t^+}|_{\Pi}$

reconstructed. The notations Ξ_k^t and \aleph_k^t are used to store the sets of invalid and unsafe transitions by Definition 4 above, given $A_{\neg}(t)$ in **Sense**_k^t and $\tilde{\Pi}_k^t$ from Section IV-B.

B. Validity and Safety

Now that we have derived the updated product automaton $\mathcal{A}_{p,k}^{t^+}$, an accepting run $R_k^{t^-}$ of $\tilde{\mathcal{A}}_{p,k}^{t^-}$ and its corresponding plan $\tau_k^{t^-}$, two natural questions arise: 1. is $\tau_k^{t^-}$ still valid or safe? 2. if not, how can we modify $\tau_k^{t^-}$ such that it remains valid and safe for $\tilde{\mathcal{T}}_k^{t^+}$ and φ_k ? Denote by $edge(R_k^{t^-})$ the set of transitions appearing in $R_k^{t^-}$. The first question is answered by the following theorem.

Theorem 3: Assume $R_k^{t^-}$ is an accepting run of $\tilde{\mathcal{A}}_{p,k}^{t^-}$. Let $\tilde{\mathcal{A}}_{p,k}^{t^-}$ be updated to $\tilde{\mathcal{A}}_{p,k}^{t^+}$. Then (i) $\tau_k^{t^-}$ remains valid if and only if $\Xi_k^t \cap \text{edge}(R_k^{t^-}) = \emptyset$; (ii) $\tau_k^{t^-}$ remains safe if $\aleph_k^t \cap \text{edge}(R_k^{t^-}) = \emptyset$.

Proof: Since $R_k^{t^-}$ is an accepting run of $\tilde{\mathcal{A}}_{p,k}^{t^-}, \tau_k^{t^-}$ is both valid and safe for $\mathcal{T}_k^{t^-}$ by Theorem 2. If $\Xi_k^t \cap edge(R_k^{t^-}) = \emptyset$ and $\aleph_k^t \cap edge(R_k^{t^-}) = \emptyset$, $R_k^{t^-}$ does not contain any invalid or unsafe transitions. Thus $R_k^{t^-}$ is still an accepting run of $\tilde{\mathcal{A}}_{p,k}^{t^+}$. Then (i) by Theorem 2, $\tau_k^{t^-}$ is still valid. On the other hand, if $\Xi_k^t \cap edge(R_k^{t^-}) \neq \emptyset$, $R_k^{t^-}|_{\Pi}$ contains at least one invalid transition, thus not valid by Definition 1; (ii) by Theorem 2, $\tau_k^{t^-}$ is also safe.

One possible solution for the second question could be to recall Algorithm 1 with respect to $\tilde{\mathcal{A}}_{p,k}^{t^+}$, which generates a new optimal accepting run of $\tilde{\mathcal{A}}_{p,k}^{t^+}$. But this method requires a complete new search of $\tilde{\mathcal{A}}_{p,k}^{t^+}$, which is computationally inefficient in case of large \mathcal{T}_k and complex φ_k , especially when $\tilde{\mathcal{A}}_{p,k}$ has to be updated frequently.

Instead we are interested in revising $R_{k}^{t^{-}}$ locally such that it fulfils the accepting condition of $\tilde{\mathcal{A}}_{p,k}^{t^{+}}$, but not necessarily optimal. The idea is that since most of the path segments belonging to $\tilde{\mathcal{A}}_{p,k}^{t^{-}}$ remain unchanged for $\tilde{\mathcal{A}}_{p,k}^{t^{+}}$, we only need to make up the edges that are invalid or unsafe in $\tilde{\mathcal{A}}_{p,k}^{t^{+}}$. In our earlier work [12], we proposed a real-time revising algorithm to revise an accepting run whenever this accepting run contains invalid transitions due to updates in the product automaton. Function Revise() from Algorithm 5 in [12] essentially determines the location of the removed transition in $R_{k}^{t^{-}}$ and locally finds a "bridging" segment that make up the removed transition by depth-first search, such that the accepting run becomes valid. The updated $R_{k}^{t^{+}}$ is projected onto II as the motion plan at time t. It is proved by Theorem 3 in [12] that an accepting run of $\tilde{\mathcal{A}}_{p,k}^{t}$ can always be found by this algorithm if there exists one. It is summarized in Algorithm 3, where Revise() is the main tool.

C. Optimality

Algorithm 3 offers a way to locally revise the invalid or unsafe plan, which however does not maintain the cost optimality as Algorithm 1. We propose the following criterion to ensure the optimality check. Denote by Υ_k^t the accumulated number of number of changes in $\tilde{\mathcal{A}}_{p,k}$: $\Upsilon_k^{t+} =$ $\Upsilon_k^{t-} + |\Xi_k^t| + |\aleph_k^t|$. Denote by Δ_k^t the timespan since the last time the optimal planner Algorithm 1 is called. Let the thresholds $N_k^{\text{call}}, T_k^{\text{call}} \ge 0$ be chosen freely by each agent. Then whenever at least one of the following condition holds: (1) $\Upsilon_k^t \ge N_k^{\text{call}}$; (2) $\Delta_k^t \ge T_k^{\text{call}}$, Algorithm 1 is called with respect the latest product automaton $\tilde{\mathcal{A}}_{p,k}^t$ to derive the optimal motion plan $\tau_{\text{opt},k}^t$, but using agent k's current position as the initial state in its FTS. Both Υ_k^{t+} and Δ_k^{t+} are reset to zero after Algorithm 1 is called.

VI. OVERALL ARCHITECTURE

The overall architecture is summarized as follows: When the system starts, each agent synthesizes its own initial motion plan using Algorithm 1. It sends requests to neighboring agents that it has not requested before. Then it checks if it receives any reply, sensing or request messages, based on which it replies to its subscribers, and updates its FTS and product automaton. At last, it decides whether the revising algorithm or the optimal planner should be called based on the triggering condition.

Theorem 4: For each agent k at any time $t \ge 0$, its motion plan τ_k^t is always valid and safe. Moreover, for any $t' \ge 0$, there exists time $t \in [t', t' + T_k^{call}]$ such that τ_k^t is the optimal motion plan for \mathcal{T}_k^t and φ_k .

Proof: It is guaranteed that an accepting run for $\tilde{\mathcal{A}}_{p,k}^{t^+}$ will be found if there exists one by both Algorithms 1 and 3. As a result, R_k^t is always an accepting run for $\tilde{\mathcal{A}}_{p,k}^t$, meaning that the corresponding motion plan τ_k^t is always valid and safe. Due to the triggering condition described earlier, Algorithm 1 is called at least once within any time period with length T_k^{call} , which completes the proof.

The implementation of the discrete motion plan involves navigating the agents from one region to the next goal region accordingly. The correctness of the proposed solutions follows from Theorem 4. Let $|\tilde{\mathcal{A}}_{p,k}|$ be the size of $\tilde{\mathcal{A}}_{p,k}$. Algorithm 1 runs in $\mathcal{O}(|\tilde{\mathcal{A}}_{p,k}| \cdot \log |\tilde{\mathcal{A}}_{p,k}| \cdot (|Q'_0| + |\mathcal{F}'|))$.

VII. CASE STUDY

In this case study, we apply the framework to a team of 9 autonomous robots: three of them are aerial vehicles and the rest are ground vehicles. All modules are implemented in Python 2.7. All simulations are carried out on a desktop computer (3.06 GHz Duo CPU and 8GB of RAM).

A. Workspace and Agent Description

As shown Fig. 1, the workspace we consider is a 9×9 square consisting of 5 base stations (in yellow, denoted by $b1 \cdots, b5$): four in the corners and one in the middle.



Fig. 1. Left: the actual workspace model as described in Section VII-A. Right: the final optimal motion plan for aerial vehicles, which corresponds to the final optimal plan in Fig. 3. It satisfies both φ_{Ari}^{hard} and φ_{Ari}^{soft} .



Fig. 2. The final optimal motion plan for agents Gw_i (left) and for agents Gf_i (right). It can be seen that water (in blue) or food (in green) is fetched before any base station and all base stations are visited infinitely often, while at the same time obstacle avoidance is ensured.

Besides, they are no-fly zone (in cyan, denoted by nfly) and numerous obstacles (in red, denoted by obs). Food (in green, denoted by food) and water (in blue, denoted by water) resources are scattered in the free space.

Three aerial vehicles (Ar_1, Ar_2, Ar_3) start randomly from one of the base stations. They have the hard specification "repetitively visit at least one of the base stations, while avoiding the no-fly zone", and soft specification "visit all base stations infinitely often". Initially, they do not know the location of base station b5, no-fly zone, nor the location of water and food resources. They have a sensing radius of 5 in the x-y coordinates. In LTL formulas, the specifications are given as $\varphi_{\mathsf{Ar}\ i}^{\mathsf{hard}} = (\Box \neg \mathsf{nfly}) \land (\Box \Diamond (\mathsf{bl} \lor \mathsf{b2} \lor \mathsf{b3} \lor \mathsf{b4} \lor \mathsf{b5}))$ and $\varphi_{Ar_{i}}^{\text{soft}} = (\Box(\Diamond b1 \land \Diamond b2 \land \Diamond b3 \land \Diamond b4 \land \Diamond b5))$, for i = 1, 2, 3. The other six robots are ground vehicles: three of them collect food (denoted by Gf 1, Gf 2, Gf 3) and the rest for water (denoted by Gw 1, Gw 2, Gw 3), to supply the base stations. The hard specification is "avoid all obstacles and repetitively collect water (or food) resources to at least one base station" and the soft specification is "supply all base stations infinitely often". Initially, they start randomly from one base station and only knows the location of one water (or food) resource, but not the obstacles and other base stations. In LTL formulas, the hard and soft specifications for Gw_i are given by $\varphi_{Gw_i}^{hard} = (\Box \Diamond \neg obs) \land$ Order and $\varphi_{Gw_i}^{soft} = \varphi_{Ar_i}^{soft}$, where i = 1, 2, 3; Order = $(\Box \Diamond water) \land (\Box (water \Rightarrow \bigcirc (\neg water U(b1 \lor b2 \lor))))$ $b3 \vee b4 \vee b5))) \land (\Box((b1 \vee b2 \vee b3 \vee b4 \vee b5)) \Rightarrow$ $\bigcirc(\neg(b1 \lor b2 \lor b3 \lor b4 \lor b5)Uwater))$. They have a sensing radius of 3 in the x-y coordinates. The soft and hard specifications $\varphi_{Gw i}^{hard}$ and $\varphi_{Gw i}^{soft}$ for Gw_i sdefined similarly



Fig. 3. It shows how the optimal plans of Ar_1, Ar_2, Ar_3 evolve under the complete, clustered and dynamic communication topologies (up to down). Each plan is labelled by the cost of its prefix and suffix. The initial plan has a total cost 3352 ([32, 332]). The final plan has a total cost 563 ([4, 56]).

by replacing proposition water with food.

Clearly, the soft specification is impossible to fulfil initially for all agents as they have no knowledge about the location of all base stations. Three different communication topologies are analysed in the simulation, including a (1) complete topology, where every pair of agents are neighbors; (2) clustered topology, where the groups of Ar_i, Gw_i, Gf_i are fully connected within each group, but only Ar_1, Gw_1, Gf_1 are neighbors between different groups; (3) dynamic topology, where the communication radius is set to 5.

B. Simulation Results

The whole scheme is implemented in the experiments. For aerial vehicles the product automaton has 240 states and 10296 edges, while the ground vehicles have a product automaton with 1920 states and 63180 edges. We choose (α, γ) to be (1000, 10) to construct the product automaton. We set N_k^{call} and T_k^{call} to be (40, 20) for aerial vehicles and (80, 15) for ground vehicles. Fig. 1 and 2 show the final optimal motion plan of groups Ar i, Gf i, Gw i. It can be seen that both soft and hard specifications are fulfilled by all agents. Fig. 3 shows how the total costs of the aerial vehicles' motion plans evolve with time under different three communication topologies. Under the complete topology, three aerial vehicles converge to the same optimal plan within 10 steps while it takes 65 steps under the clustered topology and around 160 steps under the dynamic topology. Nevertheless, it can be seen that Algorithm 1 is only called sparsely following the triggering rule in Section V-C.

VIII. CONCLUSION AND DISCUSSION

We have proposed a knowledge transfer scheme for cooperative motion planning of multi-agent systems under local LTL task specifications. The workspace is assumed to be partially known. Future work could include the consideration of dependent task specifications.

REFERENCES

- A. Bhatia, L. E. Kavraki, M. Y. Vardi. Sampling-based motion planning with temporal goals. *IEEE International Conference on Robotics and Automation(ICRA)*, 2010.
- [2] A. Bhatia, M. R. Maly, L. E. Kavraki, M. Y. Vardi. Motion planning with complex goals. *IEEE Robotics & Automation Magazine*, 18(3): 55-64, 2011.
- [3] C. Baier, J.-P Katoen. Principles of model checking. *The MIT Press*, 2008.
- [4] C. Belta, V. Isler, G. J. Pappas. Discrete abstractions for robot motion planning and control in polygonal environments. *IEEE Transactions on Robotics*, 21(5): 864-874, 2005.
- [5] C. Belta, A. Bicchi, M. Egerstedt, E. Frazzoli, E. Klavins, G. J. Pappas. Symbolic planning and control of robot motion. *IEEE Robotics and Automation Magazine*, 14: 61-71, 2007.
- [6] S. Boyd, L. Vandenberghe. *Convex Optimization*, Cambridge University Press, 2009.
- [7] E. M. Clarke, O. Grumberg, D. A. Peled. Model checking. *The MIT Press*, 1999.
- [8] X. Ding, M. Kloetzer, Y. Chen, C. Belta. Automatic deployment of robotic teams. *IEEE Robotics Automation Magazine*, 18: 75-86, 2011.
- [9] G. E. Fainekos, A. Girard, H. Kress-Gazit, G. J. Pappas. Temporal Logic Motion Planning for Dynamic Mobile Robots. *Automatica*, 45(2): 343-352, 2009.
- [10] I. F. Filippidis, D. V. Dimarogonas, K. J. Kyriakopoulos. Decentralized Multi-Agent Control from Local LTL Specifications. *IEEE Conference* on Decision and Control(CDC), 2012.
- [11] P. Gastin, D. Oddoux. Fast LTL to Büchi automaton translation. International Conference on Computer Aided Verification, 2001.
- [12] M. Guo, K. H. Johansson, D. V. Dimarogonas. Revising Motion Planning under Linear Temporal Logic Specifications in Partially Known Workspaces. *IEEE International Conference on Robotics and Automation(ICRA)*, 2013.
- [13] M. Guo, D. V. Dimarogonas. Reconfiguration in Motion Planning of Single- and Multi-agent Systems under Infeasible Local LTL Specifications. *IEEE Conference on Decision and Control(CDC)*, 2013.
- [14] S. Karaman, E. Frazzoli. Sampling-based algorithms for optimal motion planning. *International Journal of Robotics Research*, 30(7): 846-894, 2011.
- [15] S. Karaman, E. Frazzoli. Vehicle routing with linear temporal logic specifications: Applications to Multi-UAV Mission Planning. *Navigation, and Control Conference in AIAA Guidance*, 2008.
- [16] M. Kloetzer, C. Belta. Automatic deployment of distributed teams of robots from temporal logic specifications. *IEEE Transactions on Robotics*, 26(1): 48-61, 2010.
- [17] M. Kloetzer, X. C. Ding, C. Belta. Multi-robot deployment from LTL specifications with reduced communication, *IEEE Conference on Decision and Control(CDC)*, 2011
- [18] D. E. Koditschek, E. Rimon. Robot navigation functions on manifolds with boundary. Advances Appl. Math., 11:412-442, 1990.
- [19] Y. Kuwata, J. Teo, G. Fiore, S. Karaman, E. Frazzoli. Real-time motion planning with applications to autonomous urban driving. *IEEE Transactions on Control Systems*, 17(5): 1105-1118, 2009.
- [20] S. M. LaValle. Planning algorithms. Cambridge University Press, 2006.
- [21] S. C. Livingston, R. M. Murray, J. W. Burdick. Backtracking temporal logic synthesis for uncertain environments. *IEEE International Conference on Robotics and Automation(ICRA)*, 2012.
- [22] M. R. Maly, M. Lahijanian, L. E. Kavraki, H. Kress-Gazit and M. Y. Vardi. Iterative temporal motion planning for hybrid systems in partially unknown environment. *Hybrid Systems: Computation and Control (HSCC)*, 2013.
- [23] J. Tumova, L. I. Reyes-Castro, S. Karaman, E. Frazzoli, and D. Rus. Minimum-violating planning with conflicting specifications. *American Control Conference*. 2013.
- [24] J. Tumova, G. Hall, S. Karaman, E. Frazzoli, and D. Rus. Leastviolating strategy synthesis with safety rules. *In Hybrid Systems: Computation and Control (HSCC)*, 2013
- [25] T. Wongpiromsarn, U. Topcu, R. Murray. Receding horizon temporal logic planning, *IEEE Transactions on Automatic Control*, 2012.
- [26] A. Ulusoy, S. L. Smith, C. Belta. Optimal Multi-robot path planning with LTL constraints: guaranteeing correctness through synchronization. *International Symposium on Distributed Autonomous Robotic Systems*, 2012.