

## MIT Open Access Articles

### *Optimality and Robustness in Multi-Robot Path Planning with Temporal Logic Constraints*

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

**Citation:** Ulusoy, A., et al. "Optimality and Robustness in Multi-Robot Path Planning with Temporal Logic Constraints." *International Journal of Robotics Research* 32 8 (2013): 889-911.

**As Published:** 10.1177/0278364913487931

**Publisher:** SAGE Publications

**Persistent URL:** <https://hdl.handle.net/1721.1/134287>

**Version:** Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

**Terms of use:** Creative Commons Attribution-Noncommercial-Share Alike



# Optimality and Robustness in Multi-Robot Path Planning with Temporal Logic Constraints

Alphan Ulusoy<sup>†</sup> Stephen L. Smith<sup>\*</sup> Xu Chu Ding<sup>\*</sup> Calin Belta<sup>†</sup> Daniela Rus<sup>‡</sup>

## Abstract

In this paper we present a method for automatic planning of optimal paths for a group of robots that satisfy a common high level mission specification. The motion of each robot is modeled as a weighted transition system, and the mission is given as a Linear Temporal Logic (LTL) formula over a set of propositions satisfied at the regions of the environment. In addition, an optimizing proposition must repeatedly be satisfied. The goal is to minimize a cost function that captures the maximum time between successive satisfactions of the optimizing proposition while guaranteeing that the formula is satisfied.

When the robots can follow a given trajectory exactly, our method computes a set of optimal satisfying paths that minimize the cost function and satisfy the LTL formula. However, if the traveling times of the robots are uncertain, then the robots may not be able to follow a given trajectory exactly, possibly violating the LTL formula during deployment. We handle such cases by leveraging the communication capabilities of the robots to guarantee correctness during deployment and provide bounds on the deviation from the optimal values. We implement and experimentally evaluate our method for various persistent surveillance tasks in a road network environment.

## 1 Introduction

In the classical *reach-avoid* robotic path planning problem (Choset et al., 2005; LaValle, 2006), the aim is to steer a robot from a given initial position to some final position while avoiding any obstacles along the way. Many methods based on the configuration space approach (Lozano-Perez, 1983) have been proposed to find such collision-free paths. If the dimension of the configuration space permits, one can use discretized approaches that utilize various graph search algorithms (Choset et al., 2005; LaValle, 2006) or continuous methods (Rimon and Koditschek, 1992) to solve this problem. Alternatively, randomized sampling-based algorithms such as Probabilistic Road Map (PRM) (Kavraki et al., 1996) or Rapidly-Exploring Random Tree (RRT) (Kuffner and LaValle, 2000) can be used to find admissible paths. However, due to the limited scope of the problem that they address, classical path planning algorithms cannot handle more complex temporal and logic mission requirements.

Complex robotic missions need a precise as well as user-friendly language for requirement specification. In this regard, Linear Temporal Logic (LTL) provides a very attractive formalism that can capture the infinite behavior of a dynamic system in an intuitive but mathematically precise manner (Baier and Katoen, 2008). Using LTL one can easily specify complex robotic missions such as “Repeatedly visit region 1. Go to region 3 before each visit to region 1. Always avoid region 2.”. Current literature on path planning and control synthesis using LTL specifications considers finite systems, which may be abstractions of their infinite counterparts (Tabuada and Pappas, 2006; Yordanov et al., 2012). Given a finite system and an LTL

---

This work was supported in part by Office of Naval Research [grant number MURI N00014-09-1051]; Army Research Office [grant number W911NF-09-1-0088]; Air Force Office of Scientific Research [grant number YIP FA9550-09-1-020]; National Science Foundation [grant number CNS-0834260]; and by Natural Sciences and Engineering Research Council of Canada.

<sup>†</sup> Division of Systems Engineering, Boston University, Boston, MA 02215 ([alphan@bu.edu](mailto:alphan@bu.edu), [cbelta@bu.edu](mailto:cbelta@bu.edu))

<sup>\*</sup> Dept. of Electrical and Computer Engineering, University of Waterloo, Waterloo ON, N2L 3G1 Canada ([stephen.smith@uwaterloo.ca](mailto:stephen.smith@uwaterloo.ca))

<sup>\*</sup> Embedded Systems and Networks, United Technologies Research Center, East Hartford, CT 06108 ([dingx@utrc.utc.com](mailto:dingx@utrc.utc.com))

<sup>‡</sup> Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA 02139 ([rus@csail.mit.edu](mailto:rus@csail.mit.edu))

mission specification, paths and control strategies that satisfy the mission can be automatically computed for deterministic (Kloetzer and Belta, 2010; Kress-Gazit et al., 2011), non-deterministic (Kress-Gazit et al., 2007; M.Kloetzer and Belta, 2008; Thomas, 2002; Yordanov et al., 2012), and probabilistic systems (Bianco and de Alfaro, 1995; Ding et al., 2011; Kwiatkowska et al., 2002). Nevertheless, finding a path that accomplishes a mission is only part of the robotic path planning problem, as there remains the question of picking a particular path from all those paths that satisfy given specifications. In this case, one can either break the tie by making an arbitrary choice or pick the best alternative in terms of safety, speed, efficiency, or some other relevant metric.

The goal of this paper is to compute optimal paths for a group of robots subject to general LTL specifications. Our approach is motivated by persistent monitoring and pickup-delivery problems, where there is an *optimizing task* that must be repeatedly completed. We aim to compute paths that satisfy the LTL specification while minimizing the maximum time between successive completions of this *optimizing task*. Previously, we provided a method that solves this problem for a single robot (Smith et al., 2011). Then, we extended our approach to multiple robots by utilizing timed automata (Ulusoy et al., 2011), and provided improved methods that are robust to uncertainties in the speeds of robots (Ulusoy et al., 2012a,b). Moving from a single robot to multiple robots requires special care, as the model of the robotic team must capture the asynchronous motion of its members. In (Kloetzer and Belta, 2010), the authors propose a method for decentralized motion of multiple robots subject to LTL specifications. Their method, however, results in sub-optimal performance as it requires the robots to travel synchronously, blocking the execution of the mission before each transition until all robots are synchronized. The vehicle routing problem (VRP) (Toth and Vigo, 2001) and its extensions to more general classes of temporal constraints (Karaman and Frazzoli, 2008a,b) also deal with finding satisfying optimal paths for a given specification. In (Karaman and Frazzoli, 2008a), the authors consider optimal vehicle routing with metric temporal logic specifications by converting the problem to a mixed integer linear program (MILP). However, their method does not apply to the missions where robots must repeatedly complete some task, as it does not allow for specifications of the form “always eventually”. Furthermore, none of these methods are robust to timing errors that can occur during deployment, as they rely on the ability of the robots to follow generated trajectories exactly for satisfaction of the mission specification. In (Quottrup et al., 2004), the authors propose a method for synthesizing controls for a team of robots subject to a computational tree logic (CTL) formula. But, they do not consider optimizing the paths of the robots. In (Chen et al., 2012), the authors propose a method for automatic synthesis of control and communication strategies for a team of robots. However, they consider finite horizon tasks given as regular expressions as opposed to infinite horizon tasks expressed in LTL that are of our interest. Moreover, their method does not consider the costs of the generated team trajectories and thus, in general, does not provide optimal solutions. Even though the authors consider LTL as the specification language for the same problem in (Chen et al., 2011), they again do not consider optimal solutions.

The contribution of this paper is threefold. First, we provide an algorithm to capture the asynchronous motion of a group of robots. Given a team of robots modeled as weighted transition systems, this algorithm constructs a new transition system that models the joint behavior of all members as a whole. Second, we provide an algorithm to compute communication strategies for a team of robots so that we can still guarantee correctness even if the robots cannot follow generated trajectories exactly during deployment. Finally, building on these two algorithms, we present a method for generating optimal paths for a group of robots satisfying general LTL formulas. Our method is general enough to address problems involving robotic teams with different capabilities. The first case that we consider is when the members of the robotic team can follow generated paths arbitrarily closely and their models have exact timing information. One such example would be a team of robots that have accurate position information and can regulate their speeds to track moving set-points that correspond to generated paths. We address such problems with our *exact* solution that generates optimal satisfying paths. However, there might also be cases where the robots lack accurate speed control and traveling times between the regions of the environment is an unknown quantity within a given interval. If this is the case, one cannot generally guarantee satisfaction of the LTL formula without additional measures. Intuitively, if during deployment the robot speeds differ from those used for planning, then the order of events can switch, which may result in the violation of the global mission specification. For such cases we propose a *robust* solution that leverages the communication capabilities of the robots to guarantee correctness and to maintain field performance in the presence of timing errors. Paths generated using this approach are robust to uncertainties in the speeds (traveling times) of robots. In addition, we

characterize the performance of the robust paths with respect to the exact solutions. Preliminary versions of parts of our approach appeared in conference proceedings (Ulusoy et al., 2011, 2012a,b). Here, we extend these preliminary works by presenting a unified approach that can handle cases with both exact and non-deterministic traveling times. We also provide full proofs, new case studies, and experiments.

The organization of the paper is as follows. In Sec. 2, we give some preliminaries in formal methods and trace-closed languages. In Sec. 3, we formally state the optimal motion planning problem for a team of robots and give an overview of our approach. In Sec. 4, we present the parts of our approach that are common to the two cases that we consider in this paper. We present our *exact* solution in Sec. 5, which applies to the cases where the models of the robots have exact timing information and the robots can follow generated trajectories exactly. In Sec. 6, we present our *robust* solution, which applies to the cases where the traveling times of the robots are uncertain and the robots communicate to guarantee correctness during deployment and maintain field performance. In Sec. 7, we present case studies and experiments for a team of robots performing persistent data gathering missions in a road network environment. We conclude with final remarks in Section 8.

## 2 Preliminaries

In this section, we introduce the notation that we use in the rest of the paper and give some definitions. We refer the reader to (Baier and Katoen, 2008; Clarke et al., 1999; Hopcroft et al., 2007) and references therein for a more complete and rigorous treatment of these topics.

For a set  $\Pi$ , we use  $|\Pi|$ ,  $2^\Pi$ ,  $\Pi^*$ , and  $\Pi^\omega$  to denote its cardinality, power set, set of finite words, and set of infinite words, respectively. We define  $\Pi^\infty = \Pi^* \cup \Pi^\omega$  and denote the empty string by  $\emptyset$ .

**Definition 2.1 (Transition System).** A (weighted) transition system (TS) is a tuple  $\mathbf{T} := (\mathcal{Q}_T, q_T^0, \delta_T, \Pi_T, \mathcal{L}_T, w_T)$ , where

- (i)  $\mathcal{Q}_T$  is a finite set of states;
- (ii)  $q_T^0 \in \mathcal{Q}_T$  is the initial state;
- (iii)  $\delta_T \subseteq \mathcal{Q}_T \times \mathcal{Q}_T$  is the transition relation;
- (iv)  $\Pi_T$  is a finite set of atomic propositions;
- (v)  $\mathcal{L}_T : \mathcal{Q}_T \rightarrow 2^{\Pi_T}$  is a map giving the set of atomic propositions satisfied in a state;
- (vi)  $w_T : \delta_T \rightarrow \mathbb{R}_{>0}$  is a map that assigns a positive weight to each transition.

We define a *run* of  $\mathbf{T}$  as an infinite sequence of states  $r_T = q^0, q^1, \dots$  such that  $q^0 = q_T^0$ ,  $q^k \in \mathcal{Q}_T$  and  $(q^k, q^{k+1}) \in \delta_T$  for all  $k \geq 0$ . A run generates an infinite *word*  $\omega_T = \mathcal{L}(q^0), \mathcal{L}(q^1), \dots$  where  $\mathcal{L}(q^k)$  is the set of atomic propositions satisfied at state  $q^k$ . A *prefix* of a run is a finite path from an initial state to a state  $q$ . A periodic *suffix* is an infinite run originating at the state  $q$  reached by the prefix, and periodically repeating a finite path, which we call the *suffix-cycle*, originating and ending at  $q$ . A run is in *prefix-suffix form* if it consists of a prefix followed by a periodic suffix.

**Definition 2.2 (LTL Formula).** An LTL formula  $\phi$  over a set of atomic propositions  $\Pi$  is defined inductively as follows (Baier and Katoen, 2008; Clarke et al., 1999):

$$\phi := \top \mid \mathbf{p} \mid \phi \vee \phi \mid \phi \wedge \phi \mid \neg \phi \mid \mathbf{X} \phi \mid \phi \mathcal{U} \phi$$

where  $\top$  is a predicate true in each state of a system,  $\mathbf{p} \in \Pi$  is an atomic proposition,  $\neg$  (negation),  $\vee$  (disjunction) and  $\wedge$  (conjunction) are standard Boolean connectives, and  $\mathbf{X}$  and  $\mathcal{U}$  are temporal operators.

LTL formulas are interpreted over infinite words (generated by the transition system  $\mathbf{T}$  from Def. 2.1 with  $\Pi_T = \Pi$ ). Informally,  $\mathbf{X} \mathbf{p}$  states that at the next position of a word, proposition  $\mathbf{p}$  is true. Formula  $\mathbf{p}_1 \mathcal{U} \mathbf{p}_2$  states that there is a future position of the word when proposition  $\mathbf{p}_2$  is true, and proposition  $\mathbf{p}_1$  is true at least until  $\mathbf{p}_2$  is true. From these temporal operators we can construct two other temporal operators: Eventually (future),  $\mathbf{F}$  defined as  $\mathbf{F} \phi := \top \mathcal{U} \phi$ , and Always (globally),  $\mathbf{G}$ , defined as  $\mathbf{G} \phi := \neg \mathbf{F} \neg \phi$ . Formula

$\mathbf{G}\phi$  states that  $\phi$  is true at all positions of the word; formula  $\mathbf{F}\phi$  states that  $\phi$  eventually becomes true in the word. More expressivity can be achieved by combining the temporal and Boolean operators. We say a run  $r_T$  satisfies  $\phi$  if and only if the word generated by  $r_T$  satisfies  $\phi$ . An LTL formula  $\phi$  over a set  $\Pi$  can be represented by a *Büchi automaton*, which is defined next.

**Definition 2.3 (Büchi Automaton).** A Büchi automaton is a tuple  $\mathbf{B} := (\mathcal{Q}_B, \mathcal{Q}_B^0, \Pi_B, \delta_B, \mathcal{F}_B)$ , where

- (i)  $\mathcal{Q}_B$  is a finite set of states;
- (ii)  $\mathcal{Q}_B^0 \subseteq \mathcal{Q}_B$  is the set of initial states;
- (iii)  $\Pi_B$  is the input alphabet;
- (iv)  $\delta_B \subseteq \mathcal{Q}_B \times \Pi_B \times \mathcal{Q}_B$  is a non-deterministic transition relation;
- (v)  $\mathcal{F}_B \subseteq \mathcal{Q}_B$  is the set of accepting (final) states.

A run of  $\mathbf{B}$  over an input word  $\omega = \omega^0, \omega^1, \dots$  is a sequence  $r_B = q^0, q^1, \dots$ , such that  $q^0 \in \mathcal{Q}_B^0$ , and  $(q^k, \omega^k, q^{k+1}) \in \delta_B$ , for all  $k \geq 0$ . A Büchi automaton  $\mathbf{B}$  accepts a word over  $\Pi_B$  if and only if at least one of the corresponding runs intersects with  $\mathcal{F}_B$  infinitely many times. For any LTL formula  $\phi$  over a set  $\Pi$ , one can construct a Büchi automaton with input alphabet  $\Pi_B = 2^\Pi$  accepting all and only words over  $2^\Pi$  that satisfy  $\phi$ . The set of all the words accepted by a Büchi automaton  $\mathbf{B}$  is called the *language* recognized by the automaton and is denoted by  $L_B$ .

Given a set  $\Pi$ , the collection of subsets  $\Pi_i \subseteq \Pi, \forall i = 1, \dots, m$  is called a *distribution* of  $\Pi$  if  $\cup_{i=1}^m \Pi_i = \Pi$ . For a word  $\omega \in \Pi^\infty$  and a subset  $\Pi_i \subseteq \Pi$ ,  $\omega \upharpoonright_{\Pi_i}$  denotes the *projection* of  $\omega$  onto  $\Pi_i$ , which is obtained by removing all the symbols in  $\omega$  that are not in  $\Pi_i$ . For a language  $L \subseteq \Pi^\infty$  and a subset  $\Pi_i \subseteq \Pi$ ,  $L \upharpoonright_{\Pi_i}$  denotes the *projection* of  $L$  onto  $\Pi_i$ , which is the set of projections of all words in  $L$  onto  $\Pi_i$ , i.e.,  $\{\omega \upharpoonright_{\Pi_i} : \omega \in L\}$ .

**Definition 2.4 (Trace-Closed Language).** Given a distribution  $\{\Pi_1, \dots, \Pi_m\}$  of  $\Pi$  and words  $\omega, \omega' \in \Pi^\infty$ ,  $\omega'$  is *trace-equivalent* to  $\omega$ , denoted  $\omega' \sim \omega$ , iff their projections onto each one of the subsets in the given distribution are equal, i.e.,  $\omega \upharpoonright_{\Pi_i} = \omega' \upharpoonright_{\Pi_i}$  for each  $i = 1, \dots, m$ . For  $\{\Pi_1, \dots, \Pi_m\}$ , the *trace-equivalence class* of  $\omega$  is given by  $[\omega] = \{\omega' \in \Pi^\infty : \omega' \upharpoonright_{\Pi_i} = \omega \upharpoonright_{\Pi_i} \forall i = 1, \dots, m\}$ . Finally, a *trace-closed language* over  $\{\Pi_1, \dots, \Pi_m\}$  is a language  $L$  such that  $[\omega] \subseteq L, \forall \omega \in L$ .

**Remark 2.5 (Optimal-Run Algorithm (Smith et al., 2011)).** The approach that we present in this paper utilizes the OPTIMAL-RUN algorithm that we previously developed in (Smith et al., 2011). The algorithm takes as input a weighted transition system modeling the motion of a robot and an LTL formula of the form  $\phi := \varphi \wedge \mathbf{GF}\pi$ . In formula  $\phi$ ,  $\pi$  is the optimizing task that must be repeatedly satisfied and  $\varphi$  is an arbitrary LTL formula for capturing other mission requirements. The OPTIMAL-RUN algorithm outputs an optimal satisfying run that satisfies  $\phi$  and minimizes the maximum time between successive satisfying instances of  $\pi$ . We refer the interested reader to (Smith et al., 2011) for more details on the OPTIMAL-RUN algorithm.

### 3 Problem Formulation and Approach

In this section we introduce the optimal multi-robot path planning problem and motivate the need for solutions that are robust to uncertain robot speeds. Let

$$\mathcal{E} = (V, \rightarrow_{\mathcal{E}}, \Pi, \mathcal{L})$$

be a graph, where  $V$  is the set of vertices,  $\rightarrow_{\mathcal{E}} \subseteq V \times V$  is the set of edges,  $\Pi$  is a finite set of atomic propositions, and  $\mathcal{L}$  is a map giving the set of atomic propositions satisfied at a vertex. In this paper,  $\mathcal{E}$  is the quotient graph of a partitioned environment, where  $V$  is a set of labels for the regions in the partition,  $\rightarrow_{\mathcal{E}}$  is the corresponding adjacency relation. For example,  $V$  can be a set of labels for the regions and intersections for a road network and  $\rightarrow_{\mathcal{E}}$  can give their connections (see Fig. 4(a)).

Consider a team of  $m$  robots moving in an environment modeled by  $\mathcal{E}$ . The motion capabilities of robot  $i \in \{1, \dots, m\}$  are represented by a transition system  $\mathbf{T}_i = (\mathcal{Q}_i, q_i^0, \delta_i, \Pi_i, \mathcal{L}_i, w_i)$ , where  $\mathcal{Q}_i \subseteq V$ ;  $q_i^0$  is the

initial vertex of robot  $i$ ;  $\delta_i \subseteq \rightarrow_{\mathcal{E}}$  is a relation modeling the capability of robot  $i$  to move among the vertices;  $\Pi_i \subseteq \Pi$  is the set of propositions that can be satisfied by robot  $i$  and  $\{\Pi_1, \dots, \Pi_m\}$  is a distribution of  $\Pi$ ;  $\mathcal{L}_i$  is a mapping from  $\mathcal{Q}_i$  to  $2^{\Pi_i}$  showing how the propositions are satisfied at vertices;  $w_i(q, q')$  captures the time for robot  $i$  to go from vertex  $q$  to  $q'$ , which we assume to be an integer. In this model, each robot travels along the edges of the corresponding transition system  $\mathbf{T}_i$ , and spends zero time at its vertices. We assume that the robots are equipped with motion primitives that allow them to deterministically move from  $q$  to  $q'$  for each  $(q, q') \in \delta_i$ .

We consider the case where this robotic team has a mission in which some particular task must be repeatedly completed and the maximum time in between successive completions of this task must be minimized. For instance, in a persistent surveillance mission (Smith et al., 2011), the global mission could be to *keep gathering data while obeying traffic rules at all times*, and the repeating task could be *gathering data*. For this example, the robots would operate according to the mission specification while ensuring that the maximum time between successive data gatherings is minimized. Consequently, we assume that there is an optimizing proposition  $\pi \in \Pi$  corresponding to this particular repeating task and consider missions specified by LTL formulae of the form

$$\phi := \varphi \wedge \mathbf{GF}\pi, \quad (1)$$

where  $\varphi$  can be any LTL formula over  $\Pi$ , and  $\mathbf{GF}\pi$  means that the proposition  $\pi$  must be repeatedly satisfied. Our aim is to plan multi-robot paths that satisfy the mission specified by  $\phi$  and minimize the maximum time between successive satisfying instances of  $\pi$ .

To state this problem formally, we assume that each run  $r_i = q_i^0, q_i^1, \dots$  of  $\mathbf{T}_i$  (robot  $i$ ) starts at  $t = 0$  and generates a word  $\omega_i = \omega_i^0, \omega_i^1, \dots$  and a corresponding sequence of time instances  $\mathbb{T}_i := t_i^0, t_i^1, \dots$  such that  $\omega_i^k = \mathcal{L}_i(q_i^k)$  is satisfied at  $t_i^k$ . To define the behavior of the team as a whole, we interpret the sequences  $\mathbb{T}_i$  as sets, take the union  $\bigcup_{i=1}^m \mathbb{T}_i$  and order this set in an ascending order to obtain the sequence  $\mathbb{T} := t^0, t^1, \dots$ . Next, we define  $\omega_{team} = \omega_{team}^0, \omega_{team}^1, \dots$  to be the word generated by the team of robots where  $\omega_{team}^k$  is the union of all propositions satisfied at  $t^k$ . Then, we define the infinite sequence  $\mathbb{T}^\pi = \mathbb{T}^\pi(1), \mathbb{T}^\pi(2), \dots$  where  $\mathbb{T}^\pi(k)$  stands for the time instance when  $\pi$  is satisfied for the  $k^{th}$  time by the team. Finally, we define the cost function

$$J(\mathbb{T}^\pi) = \limsup_{k \rightarrow +\infty} (\mathbb{T}^\pi(k+1) - \mathbb{T}^\pi(k)). \quad (2)$$

In this paper we are particularly interested in the implementability and robustness of our solutions. Thus, we consider two cases for the traveling times given by the models of the robots: The first case that we consider is when the weight  $w_i(q, q')$  of each transition  $(q, q') \in \delta_i$  is exactly the time it takes for robot  $i$  to go from  $q$  to  $q'$  for  $i = 1, \dots, m$ . This corresponds to the case when the robots can follow any given run exactly when deployed in the environment and  $\mathbb{T}^\pi$  observed during deployment is identical to the planned  $\mathbb{T}^\pi$ . The second case that we consider is when the robots lack accurate speed control and the actual time it takes for robot  $i$  to go from  $q$  to  $q'$  is an uncertain quantity  $\tilde{w}_i(q, q')$  taking values in known intervals non-deterministically. The interval of each  $\tilde{w}_i(q, q')$  is given by  $[\underline{\rho}_i w_i(q, q'), \overline{\rho}_i w_i(q, q')]$ , where  $w_i(q, q')$  is the weight of the transition  $(q, q') \in \delta_i$ ,  $\underline{\rho}_i$  and  $\overline{\rho}_i$  are the lower and upper *deviation values* of robot  $i$ , and  $0 < \underline{\rho}_i \leq 1 \leq \overline{\rho}_i$ . In this setting, we treat the weight  $w_i(q, q')$  given by  $\mathbf{T}_i$  as a nominal value, which determines the bounds of the uncertain traveling time  $\tilde{w}_i(q, q')$  along with  $\underline{\rho}_i$  and  $\overline{\rho}_i$ . We further assume that  $\underline{\rho}_i$  and  $\overline{\rho}_i$  of each robot  $i$  are known a priori. In the following, we use  $x$  and  $\tilde{x}$  to denote the *nominal* and *actual* values of some variable  $x$ , and use the expression “*in the field*” to refer to the model with uncertain traveling times. Notice that, for the case of uncertain traveling times,  $J(\mathbb{T}^\pi)$  corresponds to the *nominal* value of the cost function, whereas  $J(\tilde{\mathbb{T}}^\pi)$  is the *actual* maximum time between any two successive satisfactions of  $\pi$  during deployment, *i.e.*,

$$J(\tilde{\mathbb{T}}^\pi) = \limsup_{k \rightarrow +\infty} (\tilde{\mathbb{T}}^\pi(k+1) - \tilde{\mathbb{T}}^\pi(k)).$$

When the robots cannot follow generated trajectories exactly, the order in which the propositions are satisfied may switch during deployment. Then, the *actual* word  $\tilde{\omega}_{team}$  generated by the robotic team during its infinite asynchronous run in the field may not be the *planned* word  $\omega_{team}$ , but a trace equivalent of  $\omega_{team}$  instead, *i.e.*,  $\tilde{\omega}_{team} \in [\omega_{team}]$ . This leads to the definition of critical words.

**Definition 3.1 (Critical Words).** *Given the language  $L_B$  of the Büchi automaton that corresponds to the LTL formula  $\phi$  over  $\Pi$ , and a distribution  $\{\Pi_1, \dots, \Pi_m\}$  of  $\Pi$ , the word  $\omega_{team}$  over  $\Pi$  is a critical word if  $\exists \tilde{\omega}_{team} \in [\omega_{team}]$  such that  $\tilde{\omega}_{team} \notin L_B$ , where  $[\omega_{team}]$  is the trace-equivalence class of  $\omega$  (Def. 2.4).*

Thus, we see that if the planned word is critical and the traveling times of the robots are non-deterministic, then we may not satisfy the specification in the field. This can be formalized by noting that the optimal runs that satisfy (1) are always in a prefix-suffix form (Smith et al., 2011), where the suffix cycle is repeated infinitely often. Using this observation and Def. 3.1 we can formally define the words that can violate the LTL formula during the deployment of a robotic team with uncertain traveling times.

**Proposition 3.2.** *If the suffix cycle of the word  $\omega_{team}$  is a critical word and the traveling times of the robots are non-deterministic, then the correctness of the motion of the robotic team during its deployment cannot be guaranteed.*

*Proof.* We denote the actual word generated by the robotic team in the field by  $\tilde{\omega}_{team}$ , whereas  $\omega_{team}$  stands for the planned word. Suppose that for each robot  $\underline{\rho}_i = 1 - \epsilon$ ,  $\overline{\rho}_i = 1 + \epsilon$ , and in the suffix cycle of  $\omega_{team}$  we have  $\alpha \subseteq \omega_{team}^k$  and  $\beta \subseteq \omega_{team}^{(k+\tau)}$ , where  $\alpha$  and  $\beta$  are the propositions generated by robots  $i$  and  $j$  at positions  $k$  and  $k + \tau$  of  $\omega_{team}$ , respectively. Further assume that  $\beta$  must not occur before  $\alpha$ , because if it does,  $\tilde{\omega}_{team}$  violates  $\phi$ . Note that we are guaranteed to find such  $\alpha$  and  $\beta$  as we assume the suffix cycle to be a critical word. In the worst-case, for  $\tilde{\omega}_{team}$  to violate  $\phi$ , we must have  $(1 + \epsilon)t^k > (1 - \epsilon)t^{k+\tau}$ , where  $t^k$  is the time at which  $\omega_{team}^k$  is satisfied. Solving for  $\epsilon$ , we get  $\epsilon > (t^{k+\tau} - t^k)/(t^k + t^{k+\tau})$ . However, as the suffix is an infinite repetition of the suffix cycle,  $\lim_{k \rightarrow \infty} (t^{k+\tau} - t^k)/(t^k + t^{k+\tau}) = 0$  and  $\phi$  is violated for any  $\epsilon > 0$ . ■

**Remark 3.3 (Worst-Case Performance in the Field under Uncertain Traveling Times).** *In addition, we can consider the performance of the team during deployment in terms of the value of the cost function (2) observed in the field. Using the same arguments presented in Prop. 3.2, it can be easily shown that the worst-case field value of (2) will be the minimum of  $(J(\tilde{\mathbb{T}}_1^\pi), \dots, J(\tilde{\mathbb{T}}_m^\pi))$ , where  $\tilde{\mathbb{T}}_i^\pi$  is the time sequence of satisfactions of  $\pi$  by robot  $i$  and  $J(\tilde{\mathbb{T}}_i^\pi)$  is the maximum duration between any two successive satisfactions of  $\pi$  by robot  $i$  in the field. This effectively means that, in the worst case, there is no benefit in executing the task with multiple robots, as at some point in the future the overall performance of the team will be limited by that of a single member.*

Prop. 3.2 shows that we cannot solely rely on the planned runs to satisfy the mission when the traveling times are uncertain and the suffix cycle of the word  $\omega_{team}$  is a critical word. Thus, for such cases, it is relevant to consider the communication capabilities of the robots as one may leverage them to guarantee correctness during deployment. We can now formulate the problem that we consider in this paper.

**Problem 3.4.** *Given an LTL formula  $\phi$  over  $\Pi$  of the form (1) and a team of  $m$  robots modeled as transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$ , possibly with uncertain traveling times characterized by deviation values  $\overline{\rho}_i, \underline{\rho}_i, i = 1, \dots, m$ ; generate individual runs and communication strategies for each robot such that  $\mathbb{T}^\pi$  minimizes the cost function (2) subject to the constraint that  $\omega_{team}$ , or  $\tilde{\omega}_{team}$  in case of uncertain traveling times, satisfies  $\phi$ .*

Since we consider LTL formulas containing  $\mathbf{GF}\pi$ , this optimization problem is always well-posed. An overview of our approach is given in Fig. 1. Notice that the exact steps we take to solve Prob. 3.4 depend on whether the traveling times of the robots are uncertain or not. Nevertheless, in both solutions, we construct the team transition system  $\mathbf{T}$  that captures the joint asynchronous motion of the robots in the environment (Sec. 4.1). We then find an optimal satisfying run on  $\mathbf{T}$  using the OPTIMAL-RUN algorithm we previously developed in (Smith et al., 2011), and we project this run back to the individual  $\mathbf{T}_i, i = 1, \dots, m$  (Sec. 4.2). In the next section, we discuss these common parts of our approach before presenting our *exact* and *robust* solutions in the sections that follow.

## 4 Modeling the Team and Finding Optimal Satisfying Runs

As given in Fig. 1, there are two operations common to both of our solutions: construction of the team transition system  $\mathbf{T}$  and finding optimal satisfying runs for individual robots. In the following, we discuss these operations.

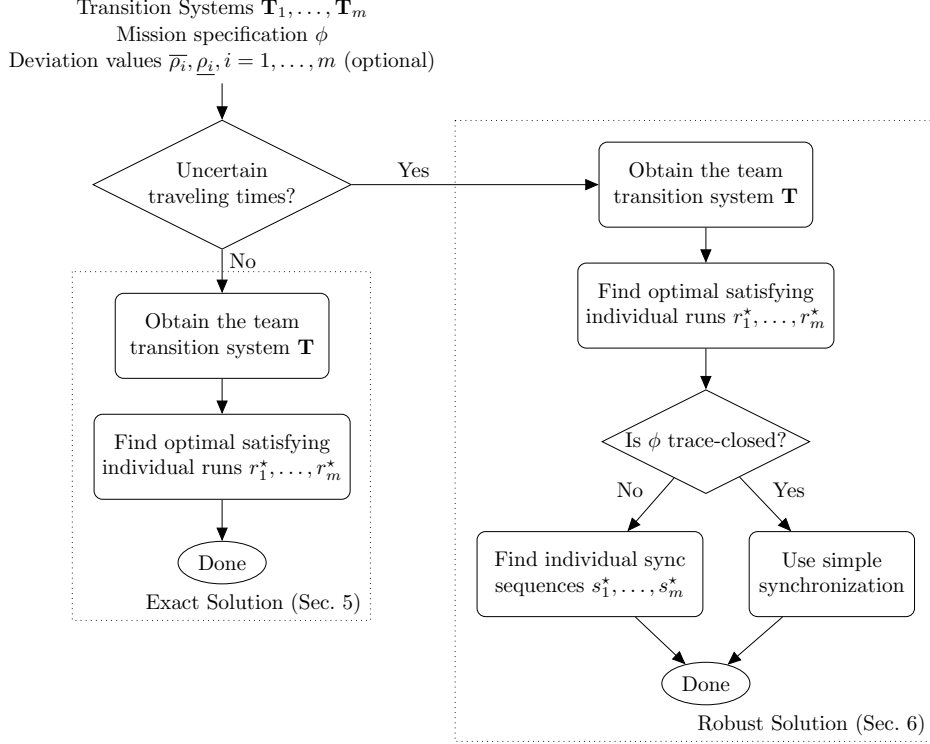


Figure 1: An overview of our approach.

## 4.1 Constructing the Team Transition System

In order to be able to optimize the motion of the team, we must capture the joint asynchronous behavior of its members as they move in the environment. Since traveling times between regions are typically not identical, we need a way to capture the states, or relative positions, of the robots regardless of whether they are at the regions in the environment or traveling in between regions. This leads to the definition of traveling states.

**Definition 4.1 (Traveling State).** *Given the transition system  $\mathbf{T}_i := (\mathcal{Q}_i, q_i^0, \delta_i, \Pi_i, \mathcal{L}_i, w_i)$  modeling robot  $i$ , we refer to a state of the form  $q_i q'_i x_i$ , where  $q_i, q'_i \in \mathcal{Q}_i$  and  $x_i > 0$ , as a traveling state, and use it to represent the instant where robot  $i$  has traveled from  $q_i$  to  $q'_i$  for  $x_i$  time units.*

To model the asynchronous motion of the team in the environment, we use a *team* transition system  $\mathbf{T} = (\mathcal{Q}_T, q_T^0, \delta_T, \Pi_T, \mathcal{L}_T, w_T)$ , where  $\mathcal{Q}_T$  is the set of states of the form  $q = (q[1], \dots, q[m])$  where  $q$  is a tuple and its  $i^{\text{th}}$  element  $q[i]$  is the state of robot  $i$ ;  $q_T^0 = (q_1^0, \dots, q_m^0)$  is the initial state of the team;  $\delta_T$  is the set of transitions;  $\Pi_T = \cup_{i=1}^m \Pi_i$  is the set of propositions;  $\mathcal{L}_T$  is a mapping from  $\mathcal{Q}_T$  to  $2^{\Pi_T}$ ;  $w_T(q, q')$  is the weight of the transition from  $q$  to  $q'$ . The states of  $\mathbf{T}$  corresponds to the instants where at least one member of the team has completed a transition on its individual transition system and is currently at a vertex while other robots may still be traveling. When robot  $i$  is at some region in the environment, we have  $q[i] \in \mathcal{Q}_i$ . If, on the other hand, robot  $i$  is traveling from  $q_i$  to  $q'_i$  and it has been  $x_i$  time units since it left  $q_i$ , we have  $q[i] = q_i q'_i x_i$ . Using this, we construct  $\mathbf{T}$  by running a depth first search on the transition systems of the individual members of the team as given in Alg. 1.

Alg. 1 is essentially a recursive depth first search (lines 4 – 17) that starts at the initial state of the team transition system  $\mathbf{T}$  (line 3). The initial state  $q_T^0$  of  $\mathbf{T}$  is defined as the tuple of the initial states of the  $m$  transition systems (line 2). Given a state  $q$  of  $\mathbf{T}$ , the function `dfsT` first generates all possible tuples of transitions that can be taken at the current states of the transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  (lines 4 – 7). The current state of transition system  $\mathbf{T}_i$  is given by the  $i^{\text{th}}$  element  $q[i]$  of the current state  $q$  of the  $\mathbf{T}$ . At line 5 of Alg. 1, we consider all possible transitions out of the current states of all transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$ . If  $q[i] \in \mathcal{Q}_i$ , *i.e.*,  $q[i]$  is a regular state of  $\mathbf{T}_i$ , then all transitions going out of this state in



---

**Algorithm 1: CONSTRUCT-TEAM-TS**

---

**Input:**  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$ .**Output:** Corresponding team transition system  $\mathbf{T}$ .1  $q_{\mathbf{T}}^0 := (q_1^0, \dots, q_m^0)$ .2  $\mathbf{dfsT}(q_{\mathbf{T}}^0)$ .3 **Function**  $\mathbf{dfsT}$ (state tuple  $q \in \mathcal{Q}_{\mathbf{T}}$ )

---

4  $q[i]$  is the  $i^{\text{th}}$  element of  $q$ .5  $t_i$  is a transition of  $\mathbf{T}_i$ , such that  $t_i \in \{(q[i], q'_i) \mid (q[i], q'_i) \in \delta_i\}$  for  $q[i] \in \mathcal{Q}_i$  and  $t_i = (q_i, q'_i)$  for  $q[i] = q_i q'_i x_i$ .6  $T := (t_1, \dots, t_m)$  is a tuple of transitions.7  $\mathcal{T}$  is the set of all such transition tuples at  $q$ .8 **foreach**  $T \in \mathcal{T}$  **do**9      $w \leftarrow$  Shortest time until a robot is at a vertex.10     Find the  $q'$  that corresponds to the new state of the team.11     **if**  $q' \notin \mathcal{Q}_{\mathbf{T}}$  **then**12         Add state  $q'$  to  $\mathcal{Q}_{\mathbf{T}}$ .13         Set  $\mathcal{L}(q') = \cup_{i=1}^m \mathcal{L}(q[i])$ .14         Add  $(q, q')$  to  $\delta_{\mathbf{T}}$  with weight  $w$ .15         Continue search from  $q'$ :  $\mathbf{dfsT}(q')$ .16     **else if**  $(q, q') \notin \delta_{\mathbf{T}}$  **then**17         Add  $(q, q')$  to  $\delta_{\mathbf{T}}$  with weight  $w$ .

---

$\mathbf{T}_i$  will be considered in the transition tuples that we will construct. Else,  $q[i]$  is a traveling state of  $\mathbf{T}_i$  of the form  $q_i q'_i x_i$ , and the only transition that can be taken is the one that is being taken, *i.e.*, the transition from  $q_i$  to  $q'_i$ . Then, we construct the set of all possible tuples of transitions that can be taken at the current states of the transition systems (lines 6–7) and process each tuple one by one (lines 8–17). In a transition tuple  $T$ , the  $i^{\text{th}}$  element gives the transition that is being taken at the current state of  $\mathbf{T}_i$ . In lines 9–10, we find the next instant where at least one transition in  $T$  has been completed and the next state  $q'$  of  $\mathbf{T}$  has been reached. The  $i^{\text{th}}$  element  $q'[i]$  of the next state  $q'$  of  $\mathbf{T}$  corresponds to the next state of  $\mathbf{T}_i$   $w$  time units after starting taking the transition  $T[i]$  at  $q[i]$ . Suppose that, the source and target states of transition  $T[i]$  are  $q_i$  and  $q'_i$ , respectively. If the transition  $T[i]$  has been completed at this point, then  $q'[i] = q'_i$ , *i.e.*, we set the next state of  $\mathbf{T}_i$  to the target state of  $T[i]$ . Otherwise,  $q'[i]$  is a traveling state of the form  $q_i q'_i x_i$  such that  $x_i = w$  if  $q[i] = q_i$ , and  $x_i = n + w$  if  $q[i] = q_i q'_i n$ . If  $q'$  is a new state (lines 11 – 15), we accordingly add it to  $\mathcal{Q}_{\mathbf{T}}$  and define its propositions. Then, we add the transition that has just been completed to  $\delta_{\mathbf{T}}$  and continue our search from this new state  $q'$ . Else, we add the transition that has just been completed to  $\delta_{\mathbf{T}}$  if required and proceed to the next transition tuple in  $\mathcal{T}$ . The algorithm concludes when all states and transitions of  $\mathbf{T}$  have been discovered.

The following proposition provides a bound on the size of the team transition system  $\mathbf{T}$ .

**Proposition 4.2.** *The number of states  $|\mathcal{Q}_{\mathbf{T}}|$  of  $\mathbf{T}$  is bounded by*

$$\prod_{i=1}^m |\mathcal{Q}_i| + (W - 1) \prod_{i=1}^m |\delta_i| \quad (3)$$

where  $W$  is the largest edge weight in all TS's.

*Proof.* The first term in (3) is the maximum number of states that we can have in the Cartesian product of  $T_i, i = 1, \dots, m$ . The second term in (3) is an upper-bound on the number of traveling states (Def. 4.1) that we can define as we construct  $\mathbf{T}$ . Here,  $\prod_{i=1}^m |\delta_i|$  is the maximum number of transitions that we can have in the Cartesian product of  $\mathbf{T}_i$ s and  $(W - 1)$  is the upper bound on the number of new traveling states per transition. Thus,  $|\mathcal{Q}_{\mathbf{T}}|$  is bounded by the sum of these two terms as given in (3). ■

**Remark 4.3 (Comparison with Naive Construction).** *One can avoid going through Alg. 1 and capture the joint behavior of the team by discretizing each transition in  $\mathbf{T}_i, i = 1, \dots, m$  to unit-length edges and taking the synchronous product of these  $m$   $\mathbf{T}_i$ 's. This approach, however, yields a much larger model whose state count is bounded by*

$$\prod_{i=1}^m \left( |\mathcal{Q}_i| + \sum_{(q, q') \in \delta_i} w_i(q, q') - |\delta_i| \right).$$

*For the case where we have  $m$  identical robots in an environment with  $Q$  vertices,  $\Delta$  edges and a largest edge weight of  $W$ , the above given bound is  $O((Q + \Delta W)^m)$ , whereas the bound given by Prop. 4.2 is  $O(Q^m + \Delta^m W)$ .*

## 4.2 Finding Optimal Satisfying Runs for Individual Robots

Once we have the transition system  $\mathbf{T}$  modeling the team, we can use the OPTIMAL-RUN algorithm (Smith et al., 2011) to obtain an optimal run  $r_{team}^*$  on  $\mathbf{T}$  that minimizes the cost function (2) and satisfies any mission specification  $\phi$  of the form (1). The optimal run  $r_{team}^*$  always consists of a finite sequence of states of  $\mathbf{T}$  (prefix), followed by infinite repetitions of another finite sequence of states of  $\mathbf{R}$  (suffix).

Given a run  $r_{team}$  of  $\mathbf{T}$ , we can finally project it onto individual robots to obtain their individual runs  $\{r_1, \dots, r_m\}$ .

**Definition 4.4 (Projection of a Run on  $\mathbf{T}$  to  $\mathbf{T}_i$ ).** *Given a run  $r_{team}$  on  $\mathbf{T}$  where  $r_{team} = q^0, q^1, \dots$ , we define its projection on  $\mathbf{T}_i$  as run  $r_i = q_i^0 q_i^1 \dots$  for all  $i = 1, \dots, m$ , such that  $q_i^k$  appears in  $r_i$  only if  $q^k[i] \in \mathcal{Q}_i$  where  $q^k[i]$  is the  $i^{th}$  element of tuple  $q^k$ .*

It can be easily seen that the set of runs  $\{r_1, \dots, r_m\}$  obtained from  $r_{team}$  using Def. 4.4 and the run  $r_{team}$  on  $\mathbf{T}$  agree with each other: The projection given in Def. 4.4 simply breaks down a sequence of tuples of states into a tuple of sequences of states, while preserving the order of the states and filtering out the traveling states. Thus, the word  $\omega$  and the time sequence  $\mathbb{T}$  generated by  $\{r_1, \dots, r_m\}$  are exactly the word  $\omega_{team}$  and the time sequence  $\mathbb{T}_{team}$  generated by  $r_{team}$ . Moreover, if the run  $r_{team}$  is in prefix-suffix form, all individual runs  $r_i$  projected from  $r_{team}$  are also in prefix-suffix form. Therefore, the individual runs projected from the optimal run  $r_{team}^*$  are always in prefix-suffix form.

## 5 Exact Solution

In this section we consider the case where the models of the robots have exact timing information and the time it takes for the robots to travel between regions during deployment is exactly the time captured in their models. Consequently, if we plan a run based on the models of the robots, the run that we will observe when the robots are deployed will be exactly the planned run in the sense that the times at which robots reach the regions in the run will be exactly as planned.

To solve Prob. 3.4 in this case, we first create a model of the motion of the team in the environment. Given the individual transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  of the robots, we use Alg. 1 to construct the team transition system  $\mathbf{T}$  that captures the joint asynchronous behavior of the robots.

**Example 5.1.** *Figs. 2(a) and 2(b) illustrate the transition systems of two robots, where  $\Pi_1 = \{p_1, \pi\}$ ,  $\Pi_2 = \{p_2, p_3, \pi\}$ , and  $\Pi = \{p_1, p_2, p_3, \pi\}$ . Using Alg. 1 we construct the team transition system  $\mathbf{T}$  (Fig. 2(c)) that captures the joint asynchronous behavior of the team in 6 states.*

Next, given an LTL mission specification  $\phi$  of the form (1), we use our previous OPTIMAL-RUN algorithm (Smith et al., 2011) to generate an optimal satisfying run  $r_{team}^*$  on the team transition system  $\mathbf{T}$ . Then, we project the optimal satisfying run  $r_{team}^*$  on  $\mathbf{T}$  onto individual transition systems using Def. 4.4 to obtain individual optimal satisfying runs  $\{r_1^*, \dots, r_m^*\}$  of the robots.

**Example 5.1 Revisited.** *Running the OPTIMAL-RUN algorithm (Smith et al., 2011) for the team transition system  $\mathbf{T}$  given in Fig. 2(c), and the formula  $\phi := \mathbf{GF}\pi$  results in the optimal run*




Figure 2: Figs. (a) and (b) show the transition systems  $\mathbf{T}_1$  and  $\mathbf{T}_2$  of two robots in an environment with three vertices. The states of the transition systems correspond to vertices  $\{a, b, c\}$  and the edges represent the motion capabilities of each robot. The weights of the edges represent the traveling times between any two vertices. Propositions  $p_1, p_2, p_3$ , and  $\pi$  are shown next to the vertices where they can be satisfied by the robots. Fig. (c) shows the team automaton capturing the joint behavior of two robots in 6 states. A state labeled  $(a, b)$  means robot 1 is at region  $a$  and robot 2 is at region  $b$ , whereas a state labeled  $(ba1, c)$  means robot 1 has traveled from  $b$  to  $a$  for 1 time unit and robot 2 is at  $c$ .

$\mathbb{T}$	0	2	3	4	6	8	10	...
$r_{team}^*$	$a, a$	$b, b$	$ba1, c$	$a, b$	$b, a$	$a, b$	$b, a$	...
$\mathcal{L}_{\mathbf{T}}(\cdot)$		$p_1, p_2, \pi$	$p_3$	$p_2, \pi$	$p_1, \pi$	$p_2, \pi$	$p_1, \pi$	...
$r_1^*$	$a$	$b$		$a$	$b$	$a$	$b$	...
$r_2^*$	$a$	$b$	$c$	$b$	$a$	$b$	$a$	...

where the first row corresponds to the times when transitions occur, the second row corresponds to the run  $r_{team}^*$ , the third row shows the propositions satisfied at each position, and the last two rows correspond to the individual runs of the robots. For this run, we see that  $(a, a), (b, b), (ba1, c)$  is the prefix and  $(a, b), (b, a)$  is the suffix and will be repeated infinite number of times. Also, the time sequence of satisfaction of  $\pi$  is  $\mathbb{T}^\pi = 2, 4, 6, 8, 10, \dots$  and the cost as defined in (2) is  $J(\mathbb{T}^\pi) = 2$ . Note that, at time  $t = 3$ , the second robot has arrived at  $c$  while the first robot is still traveling from  $b$  to  $a$ , therefore  $r_1^*$  has no state at time  $t = 3$ .

We finally summarize our exact solution in Alg. 2, and show that this algorithm indeed gives a solution to Prob. 3.4 for the case where the models of the robots have exact timing information. We analyze the overall complexity of Alg. 2 in Prop. 5.3.

---

**Algorithm 2:** EXACT-MULTI-ROBOT-OPTIMAL-RUN

---

**Input:** Transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  and an LTL specification  $\phi$  of form (1).

**Output:** A set of runs  $\{r_1^*, \dots, r_m^*\}$  that both satisfies  $\phi$  and minimizes (2).

- 1 Construct the team transition system  $\mathbf{T}$  using CONSTRUCT-TEAM-TS (Alg. 1).
  - 2 Find the optimal run  $r_{team}^*$  using OPTIMAL-RUN (Smith et al., 2011).
  - 3 Project  $r_{team}^*$  onto  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  to obtain runs  $\{r_1^*, \dots, r_m^*\}$  (Def. 4.4).
- 

**Proposition 5.2.** *Alg. 2 solves Prob. 3.4.*

*Proof.* Note that Alg. 2 combines all steps outlined in this section. Run  $r_{team}^*$  obtained from Alg. OPTIMAL-RUN both satisfies  $\phi$  and minimizes (2) among all runs of  $\mathbf{T}$  (Smith et al., 2011). As discussed in Sec. 4.2, there is a one-to-one correspondence between a set of runs  $\{r_1, \dots, r_m\}$  obtained using Def. 4.4 and a run  $r_{team}$  of  $\mathbf{T}$ . Therefore,  $\{r_1^*, \dots, r_m^*\}$  as a projection of  $r_{team}^*$  onto  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  is a solution to Prob. 3.4. ■

**Proposition 5.3.** *For the case where a group of  $m$  identical robots are expected to satisfy an LTL specification  $\phi$  in a common environment with  $Q$  vertices,  $\Delta$  edges and a largest edge weight of  $W$ , the worst-case complexity of Alg. 2 is  $O((Q^m + \Delta^m W)^3 \cdot 2^{O(|\phi|)})$ .*

*Proof.* For the above mentioned case, the worst-case size of  $\mathbf{T}$  as given in (3) is  $O(Q^m + \Delta^m W)$ . In (Smith et al., 2011), the authors give the worst-case complexity of the OPTIMAL-RUN algorithm as  $O(|T|^3 \cdot 2^{O(|\phi|)})$  where  $|T|$  is the number of states of the input transition system and  $|\phi|$  is the length of the LTL specification. Then, the worst-case complexity of Alg. 2 becomes  $O((Q^m + \Delta^m W)^3 \cdot 2^{O(|\phi|)})$ . ■

## 6 Robust Solution

In this section we consider the case where the actual traveling times of the robots observed during deployment, denoted by  $\tilde{w}_i(q, q')$ , are uncertain quantities taking values in known intervals non-deterministically. Recall from Sec. 3 that,  $\tilde{w}_i(q, q')$  lies in the interval  $[\underline{\rho}_i w_i(q, q'), \overline{\rho}_i w_i(q, q')]$ , where  $w_i(q, q')$  is the nominal value given by  $\mathbf{T}_i$ ,  $\underline{\rho}_i$  and  $\overline{\rho}_i$  are the lower and upper *deviation values* of robot  $i$ , and  $0 < \underline{\rho}_i \leq 1 \leq \overline{\rho}_i$ . Thus, when the robots execute a planned run in the field, the run observed during deployment may be different from the one planned, possibly violating the mission specification. As previously discussed in Sec. 3, our solution in this case will also comprise a communication strategy so that the satisfaction of the mission specification will be guaranteed and the deviation of the field performance from optimality will be bounded.

### 6.1 Optimal Satisfying Runs and Transition Systems with Traveling States

Given the transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  of the robots and the mission specification  $\phi$ , we first construct the team transition system  $\mathbf{T}$  using Alg. 1 to model the team. Then, we use the OPTIMAL-RUN algorithm (Smith et al., 2011) to obtain a run  $r_{team}^*$  on  $\mathbf{T}$  that satisfies  $\phi$  and minimizes the cost function (2).

**Example 6.1.** *Running the OPTIMAL-RUN algorithm (Smith et al., 2011) on  $\mathbf{T}$  given in Fig. 2(c) for the formula  $\phi = \mathbf{G}(\mathbf{p}_1 \Rightarrow \mathbf{X}(\neg \mathbf{p}_1 \mathcal{U} \mathbf{p}_3)) \wedge \mathbf{GF}\pi$  results in the optimal run*

$\mathbb{T}$	0	2	3	4	5	6	...
$r_{team}^*$	$a, a$	$b, b$	$ba1, c$	$a, b$	$ab1, c$	$b, b$	...
$\mathcal{L}_{\mathbf{T}}(\cdot)$	$\emptyset$	$\mathbf{p}_1, \mathbf{p}_2, \pi$	$\mathbf{p}_3$	$\mathbf{p}_2, \pi$	$\mathbf{p}_3$	$\mathbf{p}_1, \mathbf{p}_2, \pi$	...

where the first row shows when transitions occur, the second row corresponds to the run  $r_{team}^*$ , and the last row shows the satisfying atomic propositions. For this run,  $(a, a), (b, b)$  is the finite prefix and  $(ba1, c), (a, b), (ab1, c), (b, b)$  is the suffix cycle, which will be repeated infinite number of times. Also, the time sequence  $\mathbb{T}^\pi$  of satisfaction of  $\pi$  is  $\mathbb{T}^\pi = 2, 4, 6, 8, \dots$  and the cost as defined in (2) is  $J(\mathbb{T}^\pi) = 2$ .

Since  $\mathbf{T}$  captures the asynchronous motion of the robots, the optimal satisfying run  $r_{team}^*$  on  $\mathbf{T}$  may contain some traveling states (Def. 4.1) of the form which do not appear in the individual transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  that we started with. In our *exact* solution (Sec. 5), we pruned such states as we projected  $r_{team}^*$  onto  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  to obtain  $\{r_1^*, \dots, r_m^*\}$ . But we cannot ignore such traveling states in this case, as each one of them is a candidate synchronization point for the corresponding robot as we discuss in the following subsections. Instead, we insert those traveling states into individual transition systems so that the robots will be able to synchronize with each other at those points if needed. In the following, we use  $q^k[i]$  to denote the  $i^{th}$  element of the  $k^{th}$  state tuple in  $r_{team}^*$ , which is also the state of robot  $i$  at that position of  $r_{team}^*$ . As given in Def. 4.1, a traveling state of robot  $i$  has the form  $q_i q'_i x_i$ . First, we construct the set  $\mathcal{S} = \{(i, q^k[i]) \mid q^k[i] = q_i q'_i x_i \forall k, i\}$  of all traveling states that appear in  $r_{team}^*$ . Elements of  $\mathcal{S}$  are ordered pairs where the second element is a traveling state and the first element gives the transition system this new traveling state will be added to. Next, we construct the set  $\mathcal{T} = \{(i, (q^k[i], q^{k+1}[i]), x) \mid (q^k[i] \in \mathcal{S}) \vee (q^{k+1}[i] \in \mathcal{S}), x = w_{\mathbf{T}}(q^k, q^{k+1}) \forall k, i\}$  of all transitions that involve any of the traveling states in  $r_{team}^*$ . Elements of  $\mathcal{T}$  are triplets where the second element is a transition, the third element is the weight of this transition, and the first element shows the transition system that this new transition will be added to. Then, we add the traveling states in  $\mathcal{S}$  and the transitions in  $\mathcal{T}$  to their corresponding transition systems. Finally, using Def. 4.4, we project the run  $r_{team}^*$  onto  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  to obtain the individual runs  $r_i^*, i = 1, \dots, m$ .

**Example 1 Revisited.** *For the optimal run  $r_{team}^*$  we obtained for this example, we have  $\mathcal{S} = \{(1, ab1), (1, ba1)\}$  and  $\mathcal{T} = \{(1, (a, ab1), 1), (1, (ab1, b), 1), (1, (b, ba1), 1), (1, (ba1, a), 1)\}$ . Fig. 3 illustrates the corresponding transition systems with new traveling states and transitions highlighted in red. Then, we have runs of individual robots from Def. 4.4 as  $r_1^* = a, ab1, b, ba1, a, ab1, b, ba1, a, ab1, \dots$  and  $r_2^* = a, ab1, b, c, b, c, b, c, b, c, \dots$*

**Remark 6.2.** *For most applications, adding new states and transitions to the models of the robots may imply introducing new waypoints or motion primitives at lower levels. Since the exact way in which these model changes are accommodated at lower levels is strictly application specific, we do not discuss these details here assuming that such necessary changes can be implemented.*




Figure 3: Figs. (a) and (b) show the transition systems with new traveling states and transitions that correspond to the optimal run  $r_{team}^*$  that we compute for Exp. 6.1. In Fig. (a), the new traveling states and transitions of  $\mathbf{T}_1$  are highlighted in red.

## 6.2 Synchronization for Trace-Closed Specifications and Optimality Bounds

After obtaining individual runs of the robots, we proceed by checking if the mission specification  $\phi$  is trace-closed using an algorithm adapted from (Peled et al., 1998). We say an LTL formula  $\phi$  is trace-closed if the language  $L_B$  of the corresponding Büchi automaton is trace-closed in the sense of Def. 2.4.

**Proposition 6.3.** *If the LTL formula  $\phi$  over the set  $\Pi$  is a trace-closed formula with respect to the distribution  $\{\Pi_1, \dots, \Pi_m\}$ , then it will not be violated in the field due to uncertain traveling times.*

*Proof.* From Defs. 2.4 and 3.1, we know that if we can find a run that satisfies a trace-closed LTL formula, then the word  $\omega_{team}$  corresponding to the run will not be a critical word. We use  $\tilde{\omega}_{team}$  to denote the actual word generated by the team during deployment. Since  $\omega_{team}$  is not a critical word,  $\nexists \tilde{\omega}_{team} \in [\omega_{team}]$  such that  $\tilde{\omega}_{team} \notin L_B$ . Thus, regardless of the deviation values of the robots,  $\phi$  will not be violated in the field due to uncertain traveling times as any  $\tilde{\omega}_{team} \in [\omega_{team}]$  will also be in  $L_B$ . ■

**Corollary 6.4.** *If the LTL formula  $\phi$  over the set  $\Pi$  is not trace-closed with respect to the distribution  $\{\Pi_1, \dots, \Pi_m\}$ , then  $\phi$  may be violated during deployment due to uncertain traveling times.*

*Proof.* The proof directly follows from Prop. 6.3. ■

If  $\phi$  is not trace-closed, we cannot guarantee correctness during deployment in general as shown in Cor. 6.4. In cases where the traveling times of the robots are uncertain and  $\phi$  is not trace-closed, we compute individual synchronization sequences  $\{s_1, \dots, s_m\}$  for the robots to guarantee correctness during deployment. We discuss how we generate these synchronization sequences in greater detail in Sec. 6.3. If, on the other hand, the mission specification  $\phi$  is trace-closed, we can guarantee correctness in the field without any additional measures as shown in Prop. 6.3. Nevertheless, as given in Rem. 3.3, the field performance of the team will invariably deviate from its planned value, and in the worst-case, the field performance of the team will be limited by that of a single member. To address this issue, we propose a periodic synchronization protocol (Alg. 3). As the robots execute their infinite runs in the field, they synchronize with each other periodically at the beginning of each repetition of the suffix-cycle.

---

### Algorithm 3: TRACE-CLOSED-SYNC-RUN

---

**Input:** A run  $r_i = q_i^0, q_i^1, \dots$  of robot  $i$  in prefix-suffix form.

- 1  $q_{sync} \leftarrow$  First state in the suffix-cycle.
  - 2  $k \leftarrow 0$ .
  - 3 **while** *True* **do**
  - 4     **if** *current state is  $q_{sync}$*  **then**
  - 5         **Notify** all robots.
  - 6         **Wait** until notification messages of all robots are received.
  - 7     Make transition to  $r_i^{k+1}$ .
  - 8      $k \leftarrow k + 1$ .
- 

Using this protocol, we can define a bound on the deviation from optimality, *i.e.*, the value of the cost function (2) observed in the field, as given in the following proposition.

**Proposition 6.5.** *Suppose that each robot's deviation values are bounded by  $\underline{\rho}$  and  $\bar{\rho}$  where  $\bar{\rho} \geq \underline{\rho} > 0$  (i.e.,  $\underline{\rho}_i \geq \underline{\rho}$  and  $\bar{\rho}_i \leq \bar{\rho}$  for all robots  $i$ ), and let  $J(\mathbb{T}^\pi)$  be the cost of the planned robot paths. Then, if the robots follow the protocol given in Alg. 3 the field value of the cost satisfies*

$$\overline{J(\mathbb{T}^\pi)} \leq J(\mathbb{T}^\pi)\bar{\rho} + d_s(\bar{\rho} - \underline{\rho})$$

where  $d_s$  is the planned duration of the suffix cycle.

*Proof.* The suffix consists of an infinite number of repetitions of the suffix cycle, which we denote by  $S_c$ . As given in Alg. 3, each repetition of  $S_c$  begins with a synchronization point where all robots synchronize with each other. Let  $d_s$  be the planned duration of  $S_c$ , let  $n_s$  be the number of optimizing propositions satisfied in  $S_c$ . Let us redefine  $t = 0$  to be the time when the suffix starts, and let  $\bar{\mathbb{T}}^\pi$  be a sequence of length  $n_s$  recording the  $n_s$  times that the optimizing proposition is satisfied on the first repetition of  $S_c$ . Note that, as we consider infinite runs and as the process restarts itself at the beginning of each  $S_c$  by means of the synchronization protocol given in Alg. 3, we only need to consider the first repetition of  $S_c$ . We first define

$$\begin{aligned} \underline{T}^i &= \bar{\mathbb{T}}^\pi(i)\underline{\rho} \\ \overline{T}^i &= \bar{\mathbb{T}}^\pi(i)\bar{\rho} \\ t^w &= d_s\bar{\rho} \end{aligned}$$

where,  $\underline{T}^i$  and  $\overline{T}^i$  are the earliest and latest times that the  $i$ th optimizing proposition can be satisfied, respectively. The value  $t^w$  is the latest time that the second repetition of  $S_c$  can begin. Then, for  $0 < i \leq n_s$ , the worst-case time between satisfying the  $i^{\text{th}}$  optimizing proposition and the  $(i+1)^{\text{th}}$  optimizing proposition is

$$\tau^{i,i+1} = \begin{cases} \overline{T}^{i+1} - \underline{T}^i & \text{if } 0 < i < n_s, \\ t^w + \overline{T}^1 - \underline{T}^{n_s} & \text{if } i = n_s. \end{cases} \quad (4)$$

Next, in the planned paths, multiple robots may simultaneously satisfy the  $i$ th optimizing proposition. In the field, these satisfactions will not occur simultaneously. The maximum amount of time between the first and last of these satisfying instances for the  $i^{\text{th}}$  proposition, for  $0 < i \leq n_s$ , is

$$\tau^i = \overline{T}^i - \underline{T}^i. \quad (5)$$

Finally, using (4) and (5) we obtain the upper bound on the value of the cost function (2) that will be observed during deployment as

$$\overline{J(\mathbb{T}^\pi)} = \max\{\max_i\{\tau^{i,i+1}\}, \max_i\{\tau^i\}\}. \quad (6)$$

Substituting the definitions for  $\overline{T}^i$ ,  $\underline{T}^i$ , and  $t^w$  into (4) we obtain

$$\tau^{i,i+1} = \begin{cases} \bar{\mathbb{T}}^\pi(i+1)\bar{\rho} - \bar{\mathbb{T}}^\pi(i)\underline{\rho} & \text{if } 0 < i < n_s, \\ (d_s + \bar{\mathbb{T}}^\pi(1))\bar{\rho} - \bar{\mathbb{T}}^\pi(n_s)\underline{\rho} & \text{if } i = n_s \end{cases}$$

But, we have that  $J(\mathbb{T}^\pi) \geq \bar{\mathbb{T}}^\pi(i+1) - \bar{\mathbb{T}}^\pi(i)$ , and  $J(\mathbb{T}^\pi) \geq d_s + \bar{\mathbb{T}}^\pi(1) - \bar{\mathbb{T}}^\pi(n_s)$ . In addition,  $\bar{\mathbb{T}}^\pi(1) \leq J(\mathbb{T}^\pi)$  and  $\bar{\mathbb{T}}^\pi(i) \leq d_s$  for all  $i \in \{2, \dots, n_s\}$ . Using these expressions we obtain  $\tau^{i,i+1} \leq J(\mathbb{T}^\pi)\bar{\rho} + d_s(\bar{\rho} - \underline{\rho})$ . Similarly, we get  $\tau^i \leq d_s(\bar{\rho} - \underline{\rho})$ , and thus  $\overline{J(\mathbb{T}^\pi)} \leq J(\mathbb{T}^\pi)\bar{\rho} + d_s(\bar{\rho} - \underline{\rho})$ . ■

**Remark 6.6 (Exact Bound on  $\overline{J(\mathbb{T}^\pi)}$ ).** *In Prop. 6.5, we have provided a conservative bound for ease of presentation. However, we can also calculate an exact bound on the field value of the cost  $J(\mathbb{T}^\pi)$  using a treatment similar to the proof of Prop 6.5.*

### 6.3 Synchronization for General Specifications and Guarantee of Correctness

If the traveling times of the robots are uncertain and  $\phi$  is not trace-closed, we compute individual synchronization sequences  $\{s_1, \dots, s_m\}$  for the robots to guarantee correctness during deployment. As the robots

execute their infinite runs in the field, they synchronize with each other according to the synchronization sequences that we generate using Alg. 4. The synchronization sequence  $s_i$  of robot  $i$  is an infinite sequence of pairs of sets. The  $k^{th}$  element of  $s_i$ , denoted by  $s_i^k$ , corresponds to the  $k^{th}$  element  $q_i^k$  of  $r_i^*$ . Each  $s_i^k$  is a pair of two sets of robots:  $s_i^k = (s_{i,wait}^k, s_{i,notify}^k)$ , where  $s_{i,wait}^k$  and  $s_{i,notify}^k$  are the *wait-set* and *notify-set* of  $s_i^k$ , respectively. The *wait-set* of  $s_i^k$  is the set of robots that robot  $i$  must wait for at state  $q_i^k$  before satisfying its propositions and proceeding to the next state  $q_i^{k+1}$  in  $r_i^*$ . The *notify-set* of  $s_i^k$  is the set of robots that robot  $i$  must notify as soon as it reaches state  $q_i^k$ . As we discussed earlier in Sec. 6.1, the optimal run  $r_{team}^*$  of the team and the individual optimal runs  $r_i^*, i = 1, \dots, m$  of the robots are always in prefix-suffix form. Consequently, individual synchronization sequences  $s_i$  of the robots are also in prefix-suffix form.

---

**Algorithm 4:** SYNC-SEQ

---

**Input:** Individual runs  $\{r_1^*, \dots, r_m^*\}$ , Büchi automaton  $\mathbf{B}_{-\phi}$  of  $\neg\phi$ , and models of the robots.  
**Output:** Synchronization sequence for each robot  $\{s_1, \dots, s_m\}$ .

- 1  $\mathcal{I} = \{1, \dots, m\}$ ,  $beg$  = beginning of suffix cycle,  $end$  = end of suffix cycle.
- 2  $s_{i,wait}[k] = \mathcal{I} \setminus i$  for  $i \in \mathcal{I}$  and  $k = 1, \dots, end$ .
- 3 **foreach**  $k = 1, \dots, end$  **do**
- 4     **if**  $k \neq 1$  and  $k \neq beg$  **then**
- 5         Set  $s_{i,wait}^k = \emptyset \forall i \in \mathcal{I}$ .
- 6         Construct the transition system  $\mathbf{W}$  that generates every possible  $\tilde{\omega}_{team}$  (Alg. 6).
- 7         **if** the language of  $\mathbf{B}_{-\phi} \times \mathbf{W}$  is empty **then**
- 8             Continue to next position  $k$  in run.
- 9         **else**
- 10             Set  $s_{i,wait}^k = \mathcal{I} \setminus i \forall i \in \mathcal{I}$ .
- 11             **foreach**  $i \in \mathcal{I}$  **do**
- 12                 **foreach**  $j \in \mathcal{I} \setminus i$  **do**
- 13                     Remove  $j$  from  $s_{i,wait}^k$ .
- 14                     Construct the transition system  $\mathbf{W}$  that generates every possible  $\tilde{\omega}_{team}$  (Alg. 6).
- 15                     **if** the language of  $\mathbf{B}_{-\phi} \times \mathbf{W}$  is not empty **then**
- 16                         Add  $j$  back to  $s_{i,wait}^k$ .
- 17 Define each  $s_{i,notify}^k \forall i, k$  such that  $i \in s_{j,wait}^k \Rightarrow j \in s_{i,notify}^k \forall i \in \mathcal{I}, j \in \mathcal{I}, k = 1, \dots, end$ .
- 18 The sequence  $s_i$  is an infinite repetition of its suffix-cycle, i.e.  $s_i^{beg}, \dots, s_i^{end}$ , for each  $i \in \{1, \dots, m\}$ .

---

Alg. 4 is essentially a loop (lines 3 – 16) that computes the wait-sets for each position of the runs of the robots to guarantee correctness in the field. Initially, synchronization sequences are set so that the robots wait for each other at every position of their runs (line 2). At line 4 of Alg. 4, if  $k$  is the first position of the runs, we do not modify this initial value of  $s_{i,wait}^k$ . This ensures that all robots start executing their runs in a synchronized way. We also keep this initial value of  $s_{i,wait}^k$  if  $k$  is the beginning of the suffix cycle, so that all robots synchronize with each other globally at the beginning of each suffix cycle. This lets us define a bound on the deviation from optimality, *i.e.*, the value of the cost function (2) observed in the field, as given in Prop. 6.5. For all other positions of the runs, we try to shrink the wait-set of each  $s_i^k$  so that communication effort is minimized while we can still guarantee correctness in the field (lines 5 – 16). To this end, we first consider the case where robots do not wait for each other at this position of the run (lines 5 – 8). This is actually a heuristic based on the observation that in most missions robots synchronize only occasionally. We set all wait-sets corresponding to this position to empty sets. Then, given the runs, transition systems, deviation values, and wait-sets of the robots, we use Alg. 6 to construct the transition system  $\mathbf{W}$  that generates all possible words  $\tilde{\omega}_{team}$  that can be observed in the field due to the uncertainties in the traveling times. Next, we construct the product  $\mathbf{B}_{-\phi} \times \mathbf{W}$ , where  $\mathbf{B}_{-\phi}$  is the Büchi automaton corresponding to the negation of the LTL formula  $\phi$ . If the language of this product is empty, then the robots indeed do not need to synchronize at this position. Else, we restore the previous values of the wait-sets of this position (line 10) and consider each one of the robots in robot  $i$ 's  $k^{th}$  wait-set  $s_{i,wait}^k$  one by one (lines 11 – 16). After removing

some robot  $j$  from  $s_{i,wait}^k$ , we construct  $\mathbf{W}$  and check if the language of  $\mathbf{B}_{-\phi} \times \mathbf{W}$  is empty (lines 13 – 15). If the language of the product is empty, then robot  $i$  indeed does not need to wait for robot  $j$  at the  $k^{th}$  position of its run. Thus, we keep the new value of  $s_{i,wait}^k$ . Else, we restore  $s_{i,wait}^k$  to its previous value (line 16) and proceed with the next robot in  $s_{i,wait}^k$ . Once every robot in  $s_{i,wait}^k$  is considered, we proceed with the next robot in the team, and eventually next position of the run. Notice that, the synchronization sequences generated by Alg. 4 are free from any dead-locks as line 17 ensures that if some robot  $j$  waits for robot  $i$  at position  $k$ , then robot  $i$  notifies robot  $j$  at position  $k$ . As the synchronization sequences of the robots are in prefix-suffix form and the robots synchronize with each other globally at the beginning of each suffix-cycle (line 4), at line 18, we define the rest of each synchronization sequence as an infinite repetition of its first suffix-cycle that we have just generated. For the case where the total length of the prefix and the first suffix cycle is  $K$ , worst case complexity of Alg. 4 is  $O(m^2K(W + E))$  where  $m$  is the number of robots,  $W$  is the complexity of constructing  $\mathbf{W}$ , and  $E$  is the complexity of checking emptiness of  $\mathbf{W} \times \mathbf{B}_{-\phi}$  at each iteration. If the robots need to synchronize only occasionally, *i.e.*, if the heuristic at lines 5 – 8 succeeds most of the time, then the complexity is  $O(K(W + E))$ . The synchronization protocol that the robots follow in the field is given in Alg. 5.

---

**Algorithm 5:** SYNC-RUN

---

**Input:** The run  $r_i$  and synchronization sequence  $s_i$  of robot  $i$ .

- 1  $k \leftarrow 0$ .
- 2 **while** *True* **do**
- 3     **Notify** all robots in  $s_{i,notify}^k$ .
- 4     **Wait** until notification messages of all robots in  $s_{i,wait}^k$  are received.
- 5     Make transition to  $r_i^{k+1}$  after satisfying the propositions at  $r_i^k$ .
- 6      $k \leftarrow k + 1$ .

---

We use Alg. 6 to construct the transition system  $\mathbf{W}$  that generates all possible words that can be observed in the field for a given set of runs and synchronization sequences of the robots. We must first define some new terms before getting into the details of Alg. 6. We use the term *position* to refer to the current position of a robot in its run. If some robot  $i$  has just reached the state  $r_i^k$  in its run and satisfied the corresponding propositions after waiting for all of the robots in its wait-set  $s_{i,wait}^k$  as given in Alg.5, then the position of the robot is  $k$ . If, on the other hand, robot  $i$  has left state  $r_i^{k-1}$ , but one of the above conditions has not been satisfied yet, then the position of the robot is  $(k - 1, k)$ . A *robot-position* pair is a pair of the form  $(i, p)$  meaning that the position of robot  $i$  is  $p$  which can be either an integer or a pair of integers, as discussed above. For instance, the robot-position pair  $(i, (k - 1, k))$  means robot  $i$  is on its way from state  $r_i^{k-1}$  to state  $r_i^k$ . An *event* is a set of one or more robot-position pairs that give the new positions of the corresponding robots. In case of multiple robot-position pairs, all these changes occur simultaneously. That is, the event  $\{(i, k), (j, k)\}$  means that robots  $i$  and  $j$  have just reached position  $k$  in their runs. On the other hand, the event  $\{(i, k)\}$  means that robot  $i$  has just reached position  $k$  and gives no information about the position of robot  $j$ . Finally, an *event sequence* is a list of events that occur sequentially. Now we can begin discussing Alg. 6. The states of  $\mathbf{W}$  are tuples of positions such that the  $i^{th}$  element  $q[i]$  of some state  $q \in \mathcal{Q}_{\mathbf{W}}$  gives the current position of robot  $i$ . Consequently, at line 1 we set  $(1, \dots, 1)$  to be the initial state of  $\mathbf{W}$  as we assume that the robots start their runs synchronously (Alg. 4). Alg. 6 is essentially a loop (lines 2–12) that considers all possible sequences of events that may occur in the field. To do this, Alg. 6 relies on Alg. 8 to generate pairs of event sequences and corresponding sets of states of  $\mathbf{W}$  where those event sequences start. For an event sequence and the corresponding set of start states generated using Alg. 8, Alg. 6 adds the necessary states and transitions to  $\mathbf{W}$  starting from each possible start state (lines 3–12). Then, at line 5, we consider all events in an event sequence one by one. At lines 6–9, we compute the next state  $q'$  after the event  $e$  occurs at state  $q$ . If the position of some robot  $i$  changes due to  $e$ , then  $q'[i]$  is set to the new position given in  $e$  (line 7). Else we update the position of robot  $i$  to capture its progress. If the position of robot  $i$  is already a tuple in  $q$ , *i.e.*, if robot  $i$  is already on road, then we do not change its position in  $q'$  (line 8). Else, we update the position of robot  $i$  in  $q'$  such that it starts traveling towards the next state in its run (line 9). Next, we add the new state  $q'$  with the necessary propositions and the new transition  $(q, q')$  to



$\mathbf{W}$  as required (lines 10–11). Then, we set the current state  $q$  of  $\mathbf{W}$  to  $q'$  and switch to the next event  $e$  in the event sequence. Once we process all the events in this event sequence for all start states, we repeat the same procedure for the next event sequence. Since the runs of the robots are in prefix-suffix form, Alg. 8 is designed such that it terminates once the positions of the robots reach the end of the first suffix-cycle. Since the robots start each suffix-cycle in a synchronized way (Alg. 4), at line 14 of Alg. 6 we add a transition from all those states with no outgoing transitions to the state that corresponds to the beginning of the suffix cycle. This final step concludes the construction of  $\mathbf{W}$  by capturing the periodic structure of the runs of the robots. In order not to interrupt the flow of the paper, we present and discuss the complexity of Algs. 8 and 9, which we use to generate the event sequences discussed above, in App. B. Next, we characterize the complexity of Alg. 6.

---

**Algorithm 6:** CONSTRUCT-FIELD-WORDS-TS

---

**Input:**  $\{r_1, \dots, r_m\}$ ,  $\{s_{1,wait}, \dots, s_{m,wait}\}$ ,  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$ , and  $\bar{\rho}_i, \rho_i, i = 1, \dots, m$ .

**Output:** The field words transition system  $\mathbf{W}$  that generates all possible words that can be observed in the field.

```

1  $q_{\mathbf{W}}^0 = (1, \dots, 1)$ .
2 foreach (event_seq, start_states) generated using GENERATE-EVENT-SEQ (Alg. 8) do
3   foreach  $q_{start}$  in start_states do
4      $q = q_{start}$ .
5     foreach  $e$  in event_seq do
6       foreach  $i \in \{1, \dots, m\}$  do
7         if  $(i, k_{new}) \in e$  then  $q'[i] = k_{new}$ .
8         else if  $q[i]$  is a tuple then  $q'[i] = q[i]$ .
9         else  $q'[i] = (q[i], q[i] + 1)$ .
10      if  $q'$  is not in  $\mathbf{W}$  then Add  $q'$  to  $\mathcal{Q}_{\mathbf{W}}$  with  $\mathcal{L}_{\mathbf{W}}(q') = \cup_{(i,k) \in e} \mathcal{L}_i(r_i^k)$ .
11      if  $(q, q')$  is not in  $\mathbf{W}$  then Add  $(q, q')$  to  $\delta_{\mathbf{W}}$ .
12       $q = q'$ .
13  $q_{suffix} = (beg, \dots, beg)$  where beg corresponds to the beginning of the suffix-cycle.
14 foreach  $q \in \mathcal{Q}_{\mathbf{W}}$  such that  $\nexists (q, q') \in \delta_i$  for any  $q' \in \mathcal{Q}_{\mathbf{W}}$  do add  $(q', q_{suffix})$  to  $\delta_{\mathbf{W}}$ .
```

---

**Proposition 6.7.** *For the case where the total length of the prefix and the first suffix cycle is  $K$  and the intervals of the events of a robot do not overlap, the worst-case complexity of Alg. 6 is  $O(2^m m^{4m+3} K)$  and the maximum number of states of  $\mathbf{W}$  is  $O(m^2 K)$ .*

*Proof.* From Prop. B.1, each call to Alg. 8 takes time  $O(m^{4m})$  and we have at most  $K(2m - 1)$  event sequences with at most  $m$  events each. As each of the robots can be either at a position or traveling, each of these event sequences can have at most  $2^m$  different start states. Thus, the complexity of Alg. 6 is  $O(m^{4m} K (2m - 1) 2^m m^2)$ , which is  $O(2^m m^{4m+3} K)$ . Also, since Alg. 6 creates a new state for each new event, the maximum number of states of  $\mathbf{W}$  is  $O(m^2 K)$ . ■

**Remark 6.8.** *In Prop. 6.7 we assumed, for simplicity of presentation, that there is at most one event per robot per interval. This is typically the case where the deviation values of the robots are small enough (with respect to the length of the suffix cycle) so that the intervals of different events that belong to the same robot do not overlap. A more general complexity analysis could be performed for the case where the intervals of different events of a robot are allowed to overlap, but at the cost of increased difficulty of presentation and interpretation. We employ the same assumption in Props. 6.10, B.1, and B.2 for the same reason.*

**Example 6.1 Revisited.** *For the example we have shown throughout this section, we obtain the following individual optimal runs and synchronization sequences.*

$\mathbf{T}$	0	2	3	4	5	6	...
$r_1^*$	$a$	$b$	$ba1$	$a$	$ab1$	$b$	...
$s_1$	$(\{2\}, \{2\})$	$(\emptyset, \emptyset)$	$(\{2\}, \{2\})$	$(\emptyset, \emptyset)$	$(\emptyset, \emptyset)$	$(\emptyset, \emptyset)$	...
$\mathcal{L}_1(\cdot)$		$\mathbf{p}_1, \pi$				$\mathbf{p}_1, \pi$	...
$r_2^*$	$a$	$b$	$c$	$b$	$c$	$b$	...
$s_2$	$(\{1\}, \{1\})$	$(\emptyset, \emptyset)$	$(\{1\}, \{1\})$	$(\emptyset, \emptyset)$	$(\emptyset, \emptyset)$	$(\emptyset, \emptyset)$	...
$\mathcal{L}_2(\cdot)$		$\mathbf{p}_2, \pi$	$\mathbf{p}_3$	$\mathbf{p}_2, \pi$	$\mathbf{p}_3$	$\mathbf{p}_2, \pi$	...

In a line corresponding to a synchronization sequence  $s_i$ , first and second elements of the tuple at position  $k$  are  $s_{i,wait}^k$  and  $s_{i,notify}^k$ , respectively. The symbol  $\emptyset$  denotes an empty wait-set, or notify-set, i.e., the robot does not wait for, or notify, any other robot at that position of its run.

We finally summarize our robust solution in Alg. 7, and show that it provides a solution to Prob. 3.4. We analyze the overall complexity of Alg. 7 in Prop. 6.10.

---

**Algorithm 7:** ROBUST-MULTI-ROBOT-OPTIMAL-RUN

---

**Input:** Transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$ , corresponding deviation values and an LTL specification  $\phi$  of the form (1).

**Output:** A set of runs  $\{r_1^*, \dots, r_m^*\}$  that satisfies  $\phi$  and minimizes (2), a set of synchronization sequences  $\{s_1, \dots, s_m\}$  that guarantees correctness in the field (if applicable), and the bound on the performance of the team in the field.

- 1 Construct the team transition system  $\mathbf{T}$  using Alg. 1.
  - 2 Find an optimal run  $r_{team}^*$  on  $\mathbf{T}$  using OPTIMAL-RUN (Smith et al., 2011).
  - 3 Insert new traveling states to  $r_{team}^*$  and transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  (See. Sec. 6.1).
  - 4 Obtain individual runs  $\{r_1^*, \dots, r_m^*\}$  using Def. 4.4.
  - 5 **if**  $\phi$  is not trace-closed **then**
  - 6   Generate synchronization sequences  $\{s_1, \dots, s_m\}$  using SYNC-SEQ (Alg. 4).
  - 7 Find the bound on optimality as given in Prop. 6.5.
- 

**Proposition 6.9.** *Alg. 7 solves Prob. 3.4 when the traveling times of the robots are uncertain during deployment.*

*Proof.* Note that Alg. 7 combines all steps outlined in this section. The planned word  $\omega_{team}$  generated by the entire team satisfies  $\phi$ , and minimizes (2), as shown in (Smith et al., 2011). If the mission specification  $\phi$  is trace-closed, correctness during deployment is guaranteed by construction as given in Prop. 6.3. If  $\phi$  is not trace-closed, the synchronization sequences guarantee correctness by ensuring that the  $\tilde{\omega}_{team}$  generated in the field never violates  $\phi$  for given deviation values. Therefore, Alg. 7 solves Prob. 3.4. ■

**Proposition 6.10.** *Suppose that a group of  $m$  identical robots are expected to satisfy an LTL specification  $\phi$  in a common environment with  $Q$  vertices,  $\Delta$  edges and a largest edge weight of  $W$ . Further assume that  $K$  is the total length of the prefix and the first suffix cycle of the optimal satisfying run, and the intervals of the events of a robot do not overlap. Then, for typical cases where  $m \ll Q$ ,  $K < Q$ , the complexity of Alg. 7 is  $O((Q^m + \Delta^m W)^3 \cdot 2^{O(|\phi|)})$ .*

*Proof.* For the above mentioned case, the worst-case complexity of lines 1–4 of Alg. 7 becomes  $O((Q^m + \Delta^m W)^3 \cdot 2^{O(|\phi|)})$  from Prop. 5.3. The check for trace-closedness at line 5 can be done in time  $O(2^{O(|\phi|)} 2^{2^{O(|\phi|)}})$  (Peled et al., 1998). If this check fails, we generate synchronization sequences using Alg. 4, which runs in time  $O(m^2 K(W + E))$ . From Prop. 6.7,  $W$  is  $O(2^m m^{4m+3} K)$  and the number of states of  $\mathbf{W}$  is  $O(m^2 K)$ . Thus,  $E$  is  $O(2^{O(|\phi|)} m^2 K)$  (Baier and Katoen, 2008) and Alg. 4 becomes  $O(2^m m^{4m+5} K^2 + 2^{O(|\phi|)} m^4 K^2)$ . Notice that, the check for trace-closedness at line 5 of Alg. 7 is beneficial only if  $E < W$ . If not, we can simply assume that the result is false and proceed with the generation of the synchronization sequences using Alg. 4. Then, the overall worst-case complexity of Alg. 7 is  $O((Q^m + \Delta^m W)^3 2^{O(|\phi|)} + 2^m m^{4m+5} K^2 + 2^{O(|\phi|)} m^4 K^2)$ . For typical cases where  $m \ll Q$  and  $K < Q$ , the complexity becomes  $O((Q^m + \Delta^m W)^3 \cdot 2^{O(|\phi|)})$ . ■

## 7 Implementation and Case Studies

We implemented our algorithms in python as the LTL Optimal Multi-Agent Planner (LOMAP) package, which is publicly available online<sup>1</sup>. LOMAP uses the NetworkX graph package described in (Hagberg et al., 2008) to represent various models in our implementation and the LTL2BA software described in (Gastin and Oddoux, 2001) to convert LTL specifications to Büchi automata. LOMAP also includes an enhanced version of the OPTIMAL-RUN (Smith et al., 2011) algorithm which returns the path with the shortest suffix-cycle when there are multiple optimal paths in terms of the cost function (2). Furthermore, this new version can be executed on a computer cluster in a distributed fashion to be able to solve problems with large resource requirements. For the following case studies, we took advantage of this functionality to run LOMAP on a computing cluster consisting of five *m2.2xlarge* Amazon Elastic Compute Cloud<sup>2</sup> instances each with 34.2 GB of memory and 2.67 GHz quad-core processing power. A typical usage of our package is as follows:

- (i) The user defines the transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$  that model the robots moving in the environment in a plain text file using LOMAP’s format.
- (ii) Then, the user writes a short python script that defines the mission specification expressed in LTL in the form of (1) and calls the appropriate LOMAP function.
- (iii) Finally, the trajectory of the team and the value of the cost function are returned if the mission specification can be satisfied. Otherwise, our implementation shows an error message and quits.

In the following, we present various case studies considering persistent surveillance missions in the environment shown in Figs. 4(a) and 4(b). This environment is a road network consisting of roads, intersections, and regions for data gathering and upload. In this network, road segments are connected to each other via intersections, and the surveillance target is located in the middle, surrounded by four data gathering locations. For our case studies, we considered two Pololu m3pi robots with mbed development boards. We realized the environment using lines of black tape that correspond to the roads and intersections of the road network. The robots can navigate in the environment and can sense whether they are at an intersection or not using their infrared reflection sensors. The robots can also communicate with each other and a computer using Xbee wireless modules. In our case studies, inter-robot communication is used for synchronization of the robots, whereas computer-robot communication is used for deploying the robots according to the trajectory generated using our implementation.

The robots that we consider in our experiments have uncertain traveling times. In order to obtain their upper and lower deviation values, we measured the time it takes for both of the robots to complete the cycle “U2, 10, 11, 12, 1, 2, 21, 22, 23, 9, 10, U2” in Fig. 4(c) and recorded the maximum and minimum values among 20 trials. We chose this cycle because it tests all the motion primitives of the robots: “*left-turn*, *right-turn*, *u-turn*, and *go-straight*”. The average time for both robots to complete this cycle was approximately 17 seconds. We used this information to obtain the weights of the model given in Fig. 4(c), which were used as the *nominal* values in our computations. The maximum and minimum times for robot 1 to complete this cycle were 17.67 and 16.68 seconds, respectively. The maximum and minimum times for robot 2 were 17.56 and 16.77 seconds, respectively. Using these measurements we obtained the following deviation values:  $\bar{\rho}_1 = 1.039$ ,  $\underline{\rho}_1 = 0.981$ ,  $\bar{\rho}_2 = 1.033$ ,  $\underline{\rho}_2 = 0.986$ . In the following, we take these deviation values as  $\bar{\rho}_1 = \bar{\rho}_2 = 1.04$  and  $\underline{\rho}_1 = \underline{\rho}_2 = 0.98$  after adding a small margin of safety.

Fig. 4(c) illustrates the transition systems  $\mathbf{T}_1$  and  $\mathbf{T}_2$  that model the motion of the robots in this road network. The sets of states  $\mathcal{Q}_1$  and  $\mathcal{Q}_2$  are the sets of labels assigned to intersections and regions. The transition relations  $\delta_1$  and  $\delta_2$  give how the intersections and regions are connected and the weight maps  $w_1$  and  $w_2$  capture the time it takes for robots to take a transition. For our experiments, we assume that the transition systems  $\mathbf{T}_1$  and  $\mathbf{T}_2$  are identical except for their initial states and the sets of propositions that can be satisfied at their states. To be able to differentiate between data gatherings and uploads performed at

<sup>1</sup>LTL Optimal Multi-Agent Planner (LOMAP) Python Package is available at <http://hyness.bu.edu/lomap/>.

<sup>2</sup>Amazon EC2 is a commercial cluster computing service available at <http://aws.amazon.com/ec2/>.

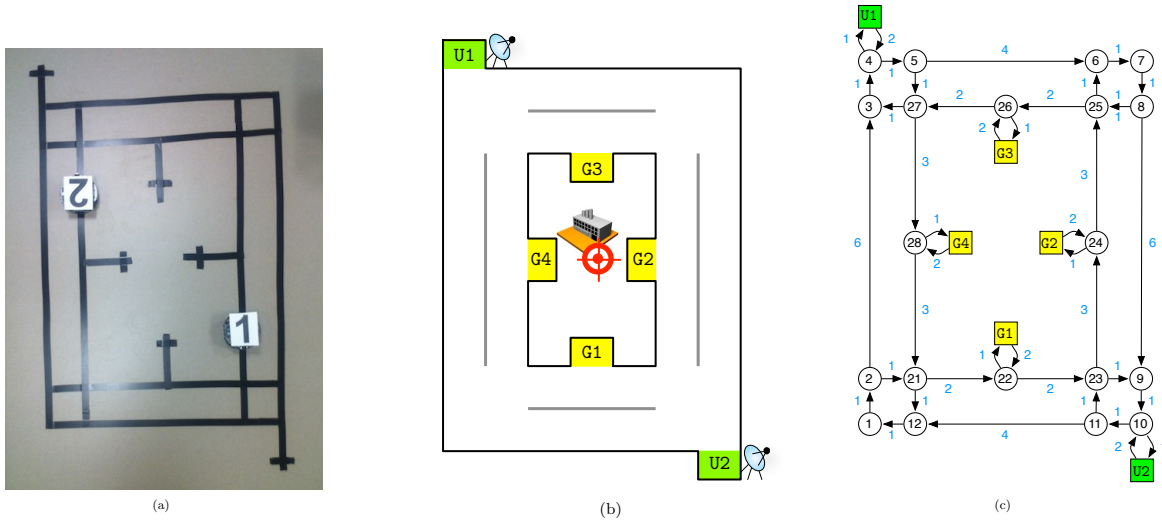


Figure 4: Fig. 4(a) shows our experimental platform where the roads are marked by black tape and the robots are labeled 1 and 2. Fig. 4(b) gives a schematic illustration of this road network. The surveillance target is in the middle. Regions highlighted in yellow are data gathering locations and regions highlighted in green are data upload locations. The transition system that models the motion of the robots is given in Fig. 4(c). The weight of each transition captures the time it takes for the robots to complete that transition.

different locations by different robots we define the set of propositions as

$$\Pi = \{\text{gather}, \text{upload}, \text{r1gather}, \text{r2gather}, \text{r1upload}, \text{r2upload}, \text{gather1}, \text{gather2}, \text{gather3}, \text{gather4}, \text{upload1}, \text{upload2}, \text{r1gather1}, \text{r1gather2}, \text{r1gather3}, \text{r1gather4}, \text{r2gather1}, \text{r2gather2}, \text{r2gather3}, \text{r2gather4}, \text{r1upload1}, \text{r1upload2}, \text{r2upload1}, \text{r2upload2}\}.$$

Propositions `gather` and `upload` mean data has been gathered and uploaded, respectively, whereas propositions of the form `gatherY` and `uploadY`, where  $Y \in \{1, 2, 3, 4\}$ , capture the locations of data gather and upload as well. For instance, `gather3` means data has been gathered at gather location 3. Propositions of the form `rXgather` and `rXupload`, where  $X \in \{1, 2\}$ , mean robot  $X$  has gathered and uploaded data, respectively. Finally, we use propositions of the form `rXgatherY` and `rXuploadY`, where  $X \in \{1, 2\}$  and  $Y \in \{1, 2, 3, 4\}$ , to capture both the location and the subject of the data gather and upload, *i.e.*, `r2Upload1` means robot 2 has uploaded data at upload location 1. Consequently, we define the sets  $\Pi_1$  and  $\Pi_2$  as

$$\begin{aligned} \Pi_1 &= \{\text{gather}, \text{upload}, \text{r1gather}, \text{r1upload}, \text{gather1}, \text{gather2}, \text{gather3}, \text{gather4}, \text{upload1}, \\ &\quad \text{upload2}, \text{r1gather1}, \text{r1gather2}, \text{r1gather3}, \text{r1gather4}, \text{r1upload1}, \text{r1upload2}\}, \text{ and} \\ \Pi_2 &= \{\text{gather}, \text{upload}, \text{r2gather}, \text{r2upload}, \text{gather1}, \text{gather2}, \text{gather3}, \text{gather4}, \text{upload1}, \\ &\quad \text{upload2}, \text{r2gather1}, \text{r2gather2}, \text{r2gather3}, \text{r2gather4}, \text{r2upload1}, \text{r2upload2}\}; \end{aligned}$$

and assign the propositions in  $\Pi_1$  and  $\Pi_2$  to the states of  $\mathbf{T}_1$  and  $\mathbf{T}_2$  as given in Tbl. 1. Note that all

Region	Propositions of Robot 1	Propositions of Robot 2
G1	{gather, gather1, r1gather, r1gather1}	{gather, gather1, r2gather, r2gather1}
G2	{gather, gather2, r1gather, r1gather2}	{gather, gather2, r2gather, r2gather2}
G3	{gather, gather3, r1gather, r1gather3}	{gather, gather3, r2gather, r2gather3}
G4	{gather, gather4, r1gather, r1gather4}	{gather, gather4, r2gather, r2gather4}
U1	{upload, upload1, r1upload, r1upload1}	{upload, upload1, r2upload, r2upload1}
U2	{upload, upload2, r1upload, r1upload2}	{upload, upload2, r2upload, r2upload2}

Table 1: Assignment of the propositions to the regions in the environment.

propositions in  $\Pi$  can be written in terms of the propositions of the last form, and therefore we could have a set  $\Pi$  consisting of eight propositions of the form  $\mathbf{rXgatherY}$  and  $\mathbf{rXuploadY}$ . However, for the sake of clarity and simplicity, we choose to define  $\Pi$  as given above, because otherwise we would have to use the long boolean expression  $\mathbf{r1Gather1} \vee \dots \vee \mathbf{r1Gather4} \vee \mathbf{r2Gather1} \vee \dots \vee \mathbf{r2Gather4}$  to express a data gather event, instead of using a single proposition, *i.e.*,  $\mathbf{gather}$ .

**Case-Study 1** – The first mission specification that we consider is as follows: “Each robot must repeatedly visit data gather locations to gather data and go to an upload location to upload their data before gathering data again. The maximum time between successive data gatherings must be minimized.” This mission specification can be expressed in LTL in the form of (1) as

$$\phi_1 := \mathbf{G}(\mathbf{r1gather} \Rightarrow \mathbf{X}(\neg \mathbf{r1gather} \mathcal{U} \mathbf{r1upload})) \wedge \\ \mathbf{G}(\mathbf{r2gather} \Rightarrow \mathbf{X}(\neg \mathbf{r2gather} \mathcal{U} \mathbf{r2upload})) \wedge \mathbf{GF}\pi,$$

where  $\pi := \mathbf{gather}$  is set as the optimizing proposition. Since the traveling times of our robots are uncertain, we use our robust solution (Sec. 5). It takes 32.5 minutes for our method to obtain an optimal satisfying team trajectory, and the cost in terms of (2) is 10. For this case, since  $\phi_1$  is trace-closed, the robots synchronize only at the beginning of their suffix-cycles. The upper bound on the value of the cost as given by Prop. 6.5 is 11.6 seconds whereas the maximum value of the cost observed in the field after 10 iterations of this trajectory was 10.66 seconds. For comparison, it also takes approximately 32.5 minutes for our exact solution to return the same trajectory with the same cost. Fig. 5(a) illustrates the optimal team trajectory that we obtain for formula  $\phi_1$ . As discussed in Sec. 4.2, optimal satisfying runs obtained using our approach always consist of a finite prefix followed by infinite repetitions of a finite suffix-cycle. In the figures that we present in this section, we omit the prefix for the sake of clarity, and use red and blue lines to illustrate the infinite periodic runs of robots 1 and 2, respectively. We use filled circles to represent the beginning of the suffix-cycles of the robots and white triangles to represent the synchronization points.

**Case-Study 2** – In some missions, sequential data gatherings at different locations may not be enough to obtain the desired information about the surveillance target. In such cases, synchronous data gatherings by multiple robots may be more desirable. For instance, one can use photographs taken synchronously from different angles to recover depth information which may be used to construct an approximate 3-d model of the surveillance target. Also, time-synchronous eavesdropping of radio communications at different locations may substantially increase the chances of recovering useful information from surveillance data. An example mission specification for such a case would be: “Robots must repeatedly gather data in a synchronous fashion, and upload their data before gathering data again.” This mission specification can be written in LTL as

$$\phi_2 := \mathbf{G}(\mathbf{gather} \Rightarrow (\mathbf{r1gather} \wedge \mathbf{r2gather})) \wedge \mathbf{G}(\mathbf{r1gather} \Rightarrow \mathbf{X}(\neg \mathbf{r1gather} \mathcal{U} \mathbf{r1upload})) \wedge \\ \mathbf{G}(\mathbf{r2gather} \Rightarrow \mathbf{X}(\neg \mathbf{r2gather} \mathcal{U} \mathbf{r2upload})) \wedge \mathbf{GF}\pi$$

where  $\pi := \mathbf{r1gather} \wedge \mathbf{r2gather}$ . Both of our robust (Sec. 6) and exact (Sec. 5) solutions take approximately 20 seconds to compute the trajectory illustrated in Fig. 5(b). The cost of this trajectory in terms of (2) is 20. The significant drop in computation from case-study 1 can be explained by the reduction in the size of the solution space in which the OPTIMAL-RUN algorithm has to work. The previous case-study requires 4656 executions of Dijkstra’s algorithm, whereas this case study requires only 680 executions of Dijkstra’s algorithm on a significantly smaller graph. We were, however, unable to execute this trajectory as our experimental setup does not allow multiple robots to be at the same region at the same time. Next, we discuss how we can address this issue and obtain a more desirable run.

**Case-Study 3** – Fig. 5(b) shows that lock-step motion of the robots is an optimal team trajectory for  $\phi_2$ . However, as our motivation for synchronous surveillance is to gather data synchronously from different locations, we can include this requirement in our specification to eliminate such undesired behaviors. Then, the mission specification can be written as

$$\phi_3 := \phi_2 \wedge \mathbf{G}(\neg(\mathbf{r1gather1} \wedge \mathbf{r2gather1}) \wedge \neg(\mathbf{r1gather2} \wedge \mathbf{r2gather2}) \wedge \\ \neg(\mathbf{r1gather3} \wedge \mathbf{r2gather3}) \wedge \neg(\mathbf{r1gather4} \wedge \mathbf{r2gather4}))$$

where  $\phi_2$  is the specification of the previous case study with  $\pi := \mathbf{r1gather} \wedge \mathbf{r2gather}$  and the rest of  $\phi_3$  forbids robots to gather data at the same place at the same time. Fig. 5(c) illustrates the optimal team

trajectory we obtain for  $\phi_3$  using our robust approach. Notice that in addition to synchronizing at the beginning of their suffix-cycles, the robots also synchronize with each other before gathering data in order not to violate the mission specification. It takes 45 seconds for our robust solution to compute this trajectory and the cost is 20. After 10 iterations of this trajectory, the maximum value of the cost observed in the field was 21 seconds, which is less than the upper bound of 22 seconds given by our approach. Extension 1 shows the execution of this trajectory by the robots.




Figure 5: Team trajectories for case studies 1, 2, and 3. Red and blue lines illustrate trajectories of robot 1 and 2, respectively. Yellow regions are data gathering locations and green regions are data upload locations. Filled circles represent the beginning of the suffix-cycles of the robots and the white triangles represent synchronization points.

**Case-Study 4** – Now, we consider the case where we need to assign each robot a specific region for data gathering while still requiring them to gather data synchronously. This is typical in scenarios where data gathering capabilities of the robots are not identical and the robots need to visit specific regions to gather useful surveillance. An example specification where robot 1 is assigned to G4 and robot 2 is assigned to G2 would be:

$$\phi_4 := \mathbf{G}(\text{gather} \Rightarrow (\text{r1gather4} \wedge \text{r2gather2})) \wedge \mathbf{G}(\text{r1gather} \Rightarrow \mathbf{X}(\neg \text{r1gather} \cup \text{r1upload})) \wedge \mathbf{G}(\text{r2gather} \Rightarrow \mathbf{X}(\neg \text{r2gather} \cup \text{r2upload})) \wedge \mathbf{GF}\pi$$

where  $\pi := \text{r1gather1} \wedge \text{r2gather4}$ . Notice that it is the sub-formula  $\mathbf{G}(\text{gather} \Rightarrow (\text{r1gather4} \wedge \text{r2gather2}))$  is  $\phi_4$  that enforces the first robot to gather data at G4 and the second robot to gather data at G2. Fig. 6(a) illustrates the optimal team trajectory we obtain for  $\phi_4$  using our robust approach. For this case, total computation time is 20 seconds and the cost is 24 with an upper bound of 26.4 seconds. After 10 iterations of this trajectory, maximum value of the cost observed in the field never exceeded 25.3 seconds.

**Case-Study 5** – In all of the case studies that we have considered so far, some of the data gathering locations have not been visited in order to optimize the team trajectory. Also, we have had the requirement that the robots must go to a dedicated upload region to upload their data before their next data gathering. However, in many cases, robots have uninterrupted links to their bases by means of some sort of wireless communication channel, and are not required to visit an upload location to upload their data. Now, we consider the case where the robots are required to visit all of the data gathering locations and are not required to visit an upload region before each data gathering. This can be expressed in LTL as

$$\phi_5 := \mathbf{GFgather1} \wedge \mathbf{GFgather2} \wedge \mathbf{GFgather3} \wedge \mathbf{GFgather4} \wedge \mathbf{GF}\pi$$

where the optimizing proposition is set as  $\pi := \text{gather}$ . Fig. 6(b) illustrates the optimal team trajectory we obtain for  $\phi_5$ . For this case, it took 23 minutes for our robust approach to obtain this trajectory. The cost

of this trajectory is 3, with an upper bound of 5.1 seconds. Since  $\phi_5$  is trace-closed, the robots synchronize only at the beginning of their suffix-cycles. It is interesting to note that the optimal solution for this case is to have robot 2 repeatedly gather data at G4 while using robot 1 to visit the remaining data gathering locations. Here, the trajectory of robot 2 minimizes the cost by gathering data as frequently as possible whereas the trajectory of robot 1 satisfies the rest of mission specification by visiting the remaining data gathering locations.

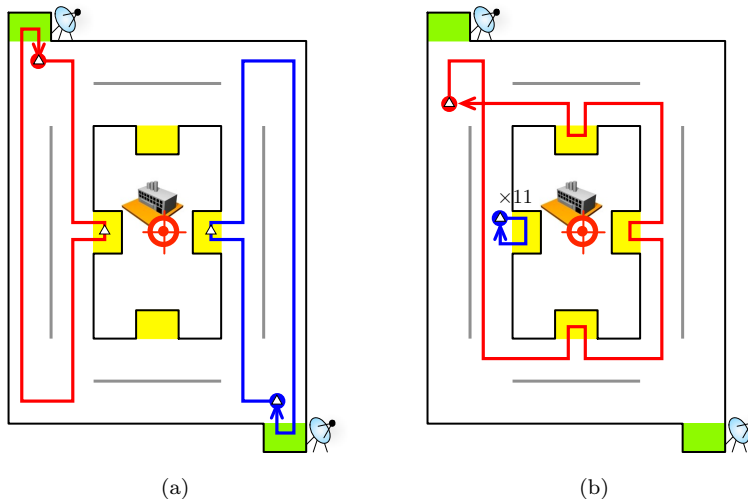


Figure 6: Team trajectories for case studies 4 and 5. Red and blue lines illustrate trajectories of robot 1 and 2, respectively. Yellow regions are data gathering locations and green regions are data upload locations. Filled circles represent the beginning of the suffix-cycles of the robots and the white triangles represent synchronization points.

## 8 Conclusions and Future Work

In this paper we presented a method for automatic planning of optimal paths for a team of robots subject to temporal logic constraints. We considered mission specifications expressed in LTL where an optimizing proposition must repeatedly be satisfied. We provided an algorithm to model the asynchronous behavior of the team as a whole, which let us extend our previous work on single robot optimal path planning to multiple robots. The motion plan that our method provides is optimal in the sense that it minimizes the maximum time in between successive satisfying instances of the optimizing proposition. Our approach is general and robust enough to handle cases where the robots cannot follow planned trajectories exactly. If the traveling times observed in the field deviate from those given by the models of the robots, our method leverages the communication capabilities of the robots to guarantee that the mission specification is never violated while overall communication effort is minimized. Our method also provides an upper bound on the difference between the performance in the field and the optimal performance in case of uncertain traveling times. We experimentally evaluated our approach and demonstrated its relevance in persistent surveillance missions in a road network environment.

In order to be able to obtain a globally optimal team trajectory, our method constructs a relatively large model that captures all members of the team and the mission specification. Thus, the main drawback of this approach is its complexity. While the method presented in this paper can be extended to Markov Decision Processes (MDPs) and different cost functions, the most rewarding direction for future research seems likely to be in the area of distributed synthesis of optimal multi robot motion plans for general mission specifications.

## Bibliography

- C. Baier and J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- A. Bianco and L. de Alfaro. Model checking of probabilistic and nondeterministic systems. In *Foundations of Software Technology and Theoretical Computer Science*, Lecture Notes in Computer Science, pages 499–513. Springer Berlin / Heidelberg, 1995.
- Y. Chen, X. C. Ding, and C. Belta. Synthesis of distributed control and communication schemes from global LTL specifications. In *2011 IEEE Conference on Decision and Control (CDC 2011)*, pages 2718–2723, Orlando, FL, 2011.
- Y. Chen, X. C. Ding, A. Stefanescu, and C. Belta. A formal approach to the deployment of distributed robotic teams. *IEEE Trans. Robotics*, 28(1):158–171, 2012.
- H. Choset, K. M. Lynch, S. Hutchinson, G. Kantor, W. Burgard, L. E. Kavraki, and S. Thrun. MIT Press, 2005.
- E. M. Clarke, D. Peled, and O. Grumberg. *Model checking*. MIT Press, 1999.
- X. C. Ding, S. L. Smith, C. Belta, and D. Rus. MDP optimal control under temporal logic constraints. In *2011 IEEE Conference on Decision and Control (CDC 2011)*, pages 532–538, Orlando, FL, 2011.
- P. Gastin and D. Oddoux. Fast LTL to Büchi automata translation. *Lecture Notes in Computer Science*, pages 53–65, 2001.
- A. A. Hagberg, D. A. Schult, and P. J. Swart. Exploring network structure, dynamics, and function using NetworkX. In *Proceedings of the 7th Python in Science Conference (SciPy2008)*, pages 11–15, Pasadena, CA, 2008.
- J. E. Hopcroft, R. Motwani, and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison Wesley, 2007.
- S. Karaman and E. Frazzoli. Vehicle routing problem with metric temporal logic specifications. In *IEEE Conf. on Decision and Control*, pages 3953–3958, Cancún, México, 2008a.
- S. Karaman and E. Frazzoli. Complex mission optimization for multiple-uavs using linear temporal logic. In *American Control Conference*, pages 2003–2009, Seattle, WA, 2008b.
- L. Kavraki, P. Svestka, J. Latombe, and M. Overmars. Probabilistic roadmaps for path planning in high-dimensional configuration spaces. *IEEE Trans. Robotics and Automation*, 12(4):566–580, 1996.
- M. Kloetzer and C. Belta. Automatic deployment of distributed teams of robots from temporal logic specifications. *IEEE Trans. Robotics*, 26(1):48–61, 2010.
- H. Kress-Gazit, G. Fainekos, and G. J. Pappas. Where’s waldo? sensor-based temporal logic motion planning. In *IEEE Intl. Conf. Robotics and Automation*, pages 3116–3121, 2007.
- H. Kress-Gazit, T. Wongpiromsarn, and U. Topcu. Correct, reactive robot control from abstraction and temporal logic specifications. *Special Issue of the IEEE Robotics & Automation Magazine on Formal Methods for Robotics and Automation*, 18:65–74, 2011.
- J. Kuffner and S. LaValle. Rrt-connect: An efficient approach to single-query path planning. In *IEEE Intl. Conf. Robotics and Automation*, page 9951001, 2000.
- M. Kwiatkowska, G. Norman, and D. Parker. Probabilistic symbolic model checking with prism: A hybrid approach. In *International Journal on Software Tools for Technology Transfer*, pages 52–66. Springer, 2002.
- S. M. LaValle. Cambridge University Press, 2006.



- T. Lozano-Perez. Spatial planning: A configuration space approach. *IEEE Trans. Comput.*, 32(2):108–120, 1983.
- M.Kloetzer and C. Belta. Dealing with non-determinism in symbolic control. In M. Egerstedt and B. Mishra, editors, *Hybrid Systems: Computation and Control: 11th International Workshop*, Lecture Notes in Computer Science, pages 287–300. Springer Berlin / Heidelberg, 2008.
- D. Peled, T. Wilke, and P. Wolper. An algorithmic approach for checking closure properties of temporal logic specifications and omega-regular languages. *Theor. Comput. Sci.*, 195(2):183–203, 1998.
- M. M. Quottrup, T. Bak, and R. Izadi-Zamanabadi. Multi-robot motion planning: A timed automata approach. In *IEEE Intl. Conf. Robotics and Automation*, pages 4417–4422, New Orleans, LA, 2004.
- E. Rimon and D. E. Koditschek. Exact robot navigation using artificial potential functions. *IEEE Trans. Robotics and Automation*, 8(5):501–518, 1992.
- S. L. Smith, J. Tůmová, C. Belta, and D. Rus. Optimal path planning for surveillance with temporal logic constraints. *Intl. Journal of Robotics Research*, 30(14):1695–1708, 2011.
- P. Tabuada and G. J. Pappas. Linear time logic control of discrete-time linear systems. *IEEE Transactions on Automatic Control*, 51(12):1862–1877, 2006.
- W. Thomas. Infinite games and verification. In *CAV*, pages 58–64, 2002.
- P. Toth and D. Vigo, editors. *The Vehicle Routing Problem*. Monographs on Discrete Mathematics and Applications. SIAM, 2001. ISBN 0898715792.
- A. Ulusoy, S. L. Smith, X. C. Ding, C. Belta, and D. Rus. Optimal path planning for surveillance with temporal logic constraints. In *IEEE/RSJ Intl. Conf. Intelligent Robots & Systems*, pages 3087–3092, San Francisco, CA, USA, Sep 2011.
- A. Ulusoy, S. L. Smith, and C. Belta. Optimal multi-robot path planning with LTL constraints: Guaranteeing correctness through synchronization. In *Intl. Symp. on Distributed and Autonomous Robotic Systems*, Baltimore, MD, USA, 2012a.
- A. Ulusoy, S. L. Smith, X. C. Ding, and C. Belta. Robust multi-robot optimal path planning with temporal logic constraints. In *IEEE Intl. Conf. Robotics and Automation*, pages 4693–4698, St. Paul, MN, USA, May 2012b.
- B. Yordanov, J. Tůmova, I. Cerna, J. Barnat, and C. Belta. Temporal logic control of discrete-time piecewise affine systems. *IEEE Trans. Automatic Control*, 57(6):1491–1504, 2012.

## A Index to Multimedia Extensions

Extension	Media Type	Description
1	Video	Execution of the trajectory in case-study 3.

## B Generation of Event Sequences

In this section, we discuss how we generate the event sequences and corresponding sets of start states that we process in Alg. 6 (Sec. 6.3). In the following, we use the terms *position*, *robot-position pair*, *event*, and *event sequence* in the same manner as we defined them in Sec. 6.3.

As discussed previously, Alg. 6 relies on Alg. 8 to construct the transition system  $\mathbf{W}$  that generates all possible words that can be observed in the field. Alg. 8 is a loop (lines 3–28) that processes a dictionary called  $tl$ , short for timeline, which we construct using Alg. 9 (line 1). We present Alg. 9 later in this section. A dictionary is a data structure that comprises a set keys, a set of values and a function that maps each key to a value. In the case of  $tl$ , the keys are time intervals and the values are sets of events. Due to non-deterministic traveling times, the time at which an event occurs in the field, in general, is not a single point but an interval. The dictionary  $tl$  captures this information by dividing the time from the beginning of the run till the end of the first suffix-cycle to disjoint intervals and by associating a set of events with each interval. The set of events that corresponds to some interval in  $tl$  is the set of events that can occur in that interval. In  $tl$ , the sets of events that correspond to different intervals are not guaranteed to be disjoint. Thus, events can span multiple intervals and can occur in either one of the intervals that they span. Suppose that the sets of events  $\{(1, 1)\}$ ,  $\{(1, 1), (2, 1)\}$ ,  $\{(1, 1)\}$  correspond to the intervals  $[0.8, 0.9]$ ,  $[0.9, 1.1]$ ,  $(1.1, 1.2]$ , respectively. Then, the event  $(1, 1)$  can occur in either one of the three intervals, whereas event  $(2, 1)$  can only occur in the interval  $[0.9, 1.1]$ .

The first part of Alg. 8 (lines 6–12) takes this fact into account while computing all possible position sequences that can be achieved by each robot at each interval. At lines 7–9, we first construct three sets of positions for each robot  $i$ : the set  $pos_{this}$  of positions that the robot can reach at this interval, the set  $pos_{prev}$  of positions that the robot can reach at either this interval or the previous interval, and the set  $pos_{next}$  of positions that the robot can reach at either this interval or the next interval. Then, at line 10, we iterate over the elements of the product  $pos_{prev} \times pos_{next}$ . For each element  $(prev, next)$  of this product set, we interpret  $prev$  as the last position that is reached in the previous interval and  $next$  as the first position that is reached in the next interval, and we obtain the remaining set of positions  $pos'_{this}$  to be reached at this interval as given in line 11. If  $pos'_{this}$  is not empty, we sort it in ascending order and add it to  $robot\_seq[i]$ , which gives the set of all possible position sequences that can be achieved by robot  $i$  at this interval.

In a given interval, events of different robots can occur in any order with respect to each other, including simultaneously. The second part of Alg. 8 (lines 14–27) addresses this by generating all possible event sequences that can be achieved by the robotic team. At line 14 we consider all combinations of position sequences that can be achieved by the robots by iterating over the elements of the product  $robot\_seq[1] \times \dots \times robot\_seq[m]$ . Next, we define  $max\_event\_cnt$  as the maximum number of events that can occur in this interval, assuming that the robots reach the positions in  $seq\_tuple$  sequentially (line 16). In order to generate all possible event sequences, we use the variable  $event\_seq$  to interpret the current interval as a box with  $max\_event\_cnt$  bins labeled  $\{1, \dots, max\_event\_cnt\}$  (line 20). For each robot  $i$ , we compute all  $len_{seq}[i]$  ordered combinations of the sequence  $\{1, \dots, max\_event\_cnt\}$  (line 17) and iterate over the elements of the product  $all\_perms[1] \times \dots \times all\_perms[m]$  (line 19). Each element of this product set is a tuple that gives how the events of individual robots are ordered with respect to the events of the other robots. Next, we obtain the event sequence corresponding to each  $perm\_tuple$  by placing the events of the robots into  $event\_seq$  according to the positions given by the  $perm\_tuple$  (lines 21–23). Notice that, as events of different robots can occur simultaneously, we may end up with some empty *bins* in  $event\_seq$ . We remove such empty entries of  $event\_seq$  at line 24. Next, we check  $event\_seq$  against the wait-sets of the robots to make sure it is a valid event sequence. To be able to do this, we need the position information of all robots, including those that do not reach a new position in  $event\_seq$ . We use the array called  $last\_pos$  to keep track of the positions of the robots. Each element  $last\_pos[i]$  of this array gives the last position reached by robot  $i$  before this interval, and we update it as given in line 28. At line 25 we use this information to check the newly computed event sequence  $event\_seq$  against the wait-sets  $\{s_{1,wait}, \dots, s_{m,wait}\}$  of the robots. If any robot

---

**Algorithm 8:** GENERATE-EVENT-SEQ

---

**Input:**  $\mathbf{W}$ ,  $\{r_1, \dots, r_m\}$ ,  $\{s_{1,wait}, \dots, s_{m,wait}\}$ ,  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$ , and  $\overline{\rho_i}, \rho_i, i = 1, \dots, m$ .

**Output:** Yields a valid event sequence and the corresponding set of starting states.

```
1 Obtain dictionary  $tl$  using COMPUTE-TIMELINE (Alg. 9).
2  $ivs$  = Sorted list of intervals of  $tl$ ,  $len_{ivs}$  = length of  $ivs$ ,  $last\_pos[i] = 0 \forall i \in \{1, \dots, m\}$ .
3 foreach  $l = 1 \dots len_{ivs}$  do
4    $events_{this} = tl[ivs[l]]$ ,  $events_{prev} = \emptyset$ ,  $events_{next} = \emptyset$ ,  $robot\_seq$  = array of  $m$  empty sets.
5   if  $l > 1$  then  $events_{prev} = tl[ivs[l-1]]$ , if  $l < len_{ivs}$  then  $events_{next} = tl[ivs[l+1]]$ .
6   foreach  $i \in \{1, \dots, m\}$  do
7      $pos_{this} = \{e | (i, e) \in events_{this}\}$ .
8      $pos_{prev} = \{e | (i, e) \in events_{prev} \cap events_{this}\}$ .
9      $pos_{next} = \{e | (i, e) \in events_{next} \cap events_{this}\}$ .
10    foreach tuple  $(prev, next)$  in  $pos_{prev} \times pos_{next}$  do
11       $pos'_{this} = \{p | p \in pos_{this}, p > prev, p < next\}$ .
12      if  $pos'_{this} \neq \emptyset$  then sort  $pos'_{this}$  in ascending order and add to  $robot\_seq[i]$ .
13  Set  $robot\_seq[i] = \{\emptyset\}$  if  $robot\_seq[i] = \emptyset \forall i \in \{1, \dots, m\}$ .
14  foreach  $seq\_tuple$  in  $robot\_seq[1] \times \dots \times robot\_seq[m]$  do
15     $len_{seq}[i] = \text{length of } seq\_tuple[i] \forall i \in \{1, \dots, m\}$ .
16     $max\_event\_cnt = \sum_{i=1}^m len_{seq}[i]$ ,  $all\_perms$  = array of  $m$  empty sets.
17     $all\_perms[i] = \text{all } len_{seq}[i] \text{ ordered combinations of } \{1, \dots, max\_event\_cnt\} \forall i \in \{1, \dots, m\}$ .
18    Set  $all\_perms[i] = \{\emptyset\}$  if  $all\_perms[i] = \emptyset \forall i \in \{1, \dots, n\}$ .
19    foreach  $perm\_tuple$  in  $all\_perms[1] \times \dots \times all\_perms[m]$  do
20       $event\_seq$  = array of  $max\_event\_cnt$  empty sets.
21      foreach  $i \in \{1, \dots, m\}$  do
22        foreach  $n \in \{1, \dots, len_{seq}[i]\}$  do
23          Add event  $(i, seq\_tuple[i][n])$  to  $event\_seq[perm\_tuple[i][n]]$ .
24      Remove those entries with  $event\_seq[i] = \emptyset$  for  $i \in \{1, \dots, max\_event\_cnt\}$ .
25      if wait-set checks fail then continue.
26       $start\_states = \{q | q \in \mathcal{Q}_W, \text{ for all } i \in \{1, \dots, m\} \text{ if } i \notin r \text{ then } q[i] = (last\_pos[i], last\_pos[i] + 1),$ 
27      else  $q[i]$  can be both  $seq\_tuple[i][1] - 1$  and  $(seq\_tuple[i][1] - 1, seq\_tuple[i][1])\}$ .
28      Yield  $(event\_seq, start\_states)$ .
29  foreach event  $(i, p)$  in  $events_{this} \setminus events_{next}$  do if  $last\_pos[i] < p$  then  $last\_pos[i] = p$ .
```

---

leaves some position earlier than any of the robots it has to wait there, then we ignore this *event\_seq* and proceed with the next one. Next, at line 26, we compute the set of start states of  $\mathbf{W}$  at which *event\_seq* can start occurring. Finally, at line 27 we yield the *event\_seq* along with the corresponding set of start states. At the next call, Alg. 8 continues execution from line 14 with the next *seq\_tuple*, and eventually from line 6 with the next interval. Once all the intervals of *tl* are considered, Alg. 8 terminates causing the loop that it is called in Alg. 6 to terminate as well.

**Proposition B.1.** *For the case where the timeline  $tl$  can be constructed in  $O(T)$  time, contains  $I$  intervals with at most  $m$  events per interval and one event per robot per interval, the worst-case per-call (lines 6–28) complexity of Alg. 8 is  $O(m^{4m})$ , and the worst-case overall complexity of Alg. 8 is  $O(I m^{4m} + T)$ .*

*Proof.* The complexity of the first part of the algorithm (lines 6–12) is  $O(m((m-1)/2)^2)$ , *i.e.*,  $O(m^3/4)$ , where we take  $pos_{this} = 1, pos_{prev} = pos_{next} = (m-1)/2$  to achieve the highest number of iterations at lines 10–12. Since we assume that there is at most one event per robot per interval, the maximum values of  $len_{seq}[i]$  and  $max\_event\_cnt$  are 1 and  $m$ . Then,  $|all\_perms[i]|$  is at most  $m$ , and the complexity of the inner loop at lines 19–27 becomes  $O(m^{m+1})$ . Since each  $|robot\_seq[i]|$  is  $O(m^3/4)$ , the complexity of the second part of the algorithm (lines 14–28) becomes  $O(m^{4m+1}/4^m)$ . As  $O(m^{4m+1}/4^m) > O(m^3/4)$  and  $m < 4^m$  for  $m \geq 1$  the per-call complexity of Alg. 8 is  $O(m^{4m})$ . Since we assume that *tl* has  $I$  intervals and the complexity of constructing *tl* is  $O(T)$ , the overall complexity of Alg. 8 becomes  $O(I m^{4m} + T)$ . ■

We use Alg. 9 to construct the dictionary *tl*, short for timeline, that we use in Alg. 8. As discussed earlier, since the runs of the robots are periodic, we consider only the prefix and the first suffix-cycle of the runs of the robots during the construction of *tl*. The first part of Alg. 9 (lines 1–7) computes the intervals in which the events of the robots can occur, *i.e.*, the intervals in which the robots can depart their current states. The interval in which some event  $(i, k)$  of robot  $i$  can occur is determined by the deviation values  $\bar{\rho}_i, \underline{\rho}_i$ , the time  $w(r_i^{k-1}, r_i^k)$  it takes for the robot to reach  $r_i^k$  from its previous position  $r_i^{k-1}$ , wait-set  $s_{i,wait}^k$  of the robot for position  $k$ , and the interval in which the robot has departed from its previous position, *i.e.*, the interval corresponding to the previous event  $(i, k-1)$ . As the robots start their runs in a synchronized way, we set the interval of the first events of all robots to  $[0, 0]$  at line 3. For all other positions, we first construct the set *waits\_for* of robots that robot  $i$  has to wait for either directly or indirectly at that position (line 4). The set of robots that robot  $i$  has to wait for directly at position  $k$  is simply given by  $s_{i,wait}^k$ . The set of robots that robot  $i$  has to wait for indirectly at position  $k$ , on the other hand, is computed recursively by forming a set equal to  $s_{i,wait}^k$ , and then, by adding the wait-sets of each robot in that set until the set stops changing. For the sake of conciseness, we also include the robot  $i$  itself in the set *waits\_for*. Next, at lines 5–6 we calculate the earliest and latest time that the event  $(i, k)$  can occur by using the models of the robots in the set *waits\_for* and the intervals of their previous events. Then, at line 7, we save the interval of event  $(i, k)$  in the *event\_ivs* array.

The second part of Alg. 9 (lines 8–28), projects the intervals in *event\_ivs* to a common *event timeline* by considering each position  $k$  of each robot  $i$ . The variable *tl* is a dictionary of sets of events keyed by intervals. To be able to use this dictionary by iterating over its keys as discussed earlier, we need to make sure that its keys, which are intervals, do not intersect with each other. To this end, we maintain the queue *projection\_queue* to hold the remaining parts of the intersecting intervals that we may need to break up during the projection. We start the projection by adding the interval of the event  $(i, k)$  to the projection queue. Then, for each interval *new\_iv* in the projection queue, we check all the intervals in *tl* to see if any of them intersects with *new\_iv*. If not, we add this interval *new\_iv* to the timeline along with its event set (line 28). If, on the other hand, the interval *new\_iv* intersects with some interval *old\_iv* in *tl*, we set the interval *int\_iv* to be the intersection of *new\_iv* and *old\_iv* and add it to the timeline with the appropriate set of events (line 17). Next, at lines 18–27 we check to see if we need to break the *old\_iv* or *new\_iv*. If *old\_iv* extends beyond *new\_iv* from the beginning or the end, we break it appropriately by defining a new entry for the extending parts and removing the old entry that corresponds to *old\_iv* from *tl*. If, on the other hand, *new\_iv* extends beyond *old\_iv*, we do not add the extending parts to *tl* directly as they may intersect with other intervals already in *tl*. Instead, we add the extending parts of *new\_iv* to the projection queue so that they are processed in the coming iterations. Alg. 9 terminates once it processes all events of all robots up to the end of the first suffix-cycle of their runs.

---

**Algorithm 9:** COMPUTE-TIMELINE

---

**Input:** Individual runs  $\{r_1, \dots, r_m\}$ , wait-sets  $\{s_{1,wait}, \dots, s_{m,wait}\}$ , transition systems  $\{\mathbf{T}_1, \dots, \mathbf{T}_m\}$ , and deviation values  $\bar{\rho}_i, \rho_i, i = 1, \dots, m$  of the robots.

**Output:** The dictionary  $tl$  of sets of events keyed by disjoint intervals.

```
1 for  $k = 1, \dots, end$  do
2   for  $i = 1, \dots, m$  do
3     if  $k$  is 1 then  $event\_ivs[i][k] = [0, 0]$  else
4        $waits\_for = \{i\} \cup \{j \mid i \text{ waits for } j \text{ at pos } k \text{ either directly or indirectly}\}.$ 
5        $earliest = \max_{j \in waits\_for} (event\_ivs[j][k-1].start + \rho_j * w_j(r_j^{k-1}, r_j^k)).$ 
6        $latest = \max_{j \in waits\_for} (event\_ivs[j][k-1].end + \bar{\rho}_j * w_j(r_j^{k-1}, r_j^k)).$ 
7        $event\_ivs[i][k] = [earliest, latest].$ 
8 for  $k = 1, \dots, end$  do
9   for  $i = 1, \dots, m$  do
10     $projection\_queue = \{event\_ivs[i][k]\}.$ 
11    foreach  $new\_iv \in projection\_queue$  do
12       $intersected = False.$ 
13      foreach  $old\_iv \in tl$  do
14         $int\_iv$  is the intersection of  $new\_iv$  and  $old\_iv.$ 
15        if  $new\_iv$  intersects with  $old\_iv$  then
16           $intersected = True.$ 
17           $tl[int\_iv] = tl[old\_iv] \cup \{(i, k)\}.$ 
18          if  $old\_iv.start < new\_iv.start$  then
19             $tl[[old\_iv.start, new\_iv.start]] = tl[old\_iv].$ 
20            Remove  $old\_iv$  from  $tl.$ 
21          if  $old\_iv.end > new\_iv.end$  then
22             $tl[(new\_iv.end, old\_iv.end)] = tl[old\_iv].$ 
23            Remove  $old\_iv$  from  $tl.$ 
24          if  $new\_iv.start < old\_iv.start$  then
25            Add  $[new\_iv.start, old\_iv.start)$  to  $projection\_queue.$ 
26          if  $new\_iv.end > old\_iv.end$  then
27            Add  $(old\_iv.end, new\_iv.end]$  to  $projection\_queue.$ 
28    if  $intersected$  is  $False$  then  $tl[new\_iv] = \{(i, k)\}$ 
29 Return  $tl.$ 
```

---

**Proposition B.2.** *For the case where the total length of the prefix and the first suffix cycle is  $K$  and the intervals of the events of a robot do not overlap, the worst-case complexity of Alg. 9 is  $O((mK)^2)$ .*

*Proof.* In the worst-case, each robot waits for every other robot, thus computation of each  $event\_ivs[i][k]$  at lines 4–7 takes time  $O(3m)$ . Then, the complexity of the first part of the algorithm (lines 1–7) is  $O(m^2K)$ . In the second part of the algorithm (lines 8–28), in the worst-case, each projected interval intersects with previously defined intervals resulting in 2 additional intervals per projection. Thus, we have at most  $2m - 1$  intervals for each position and at most  $K(2m - 1)$  intervals in total. Consequently, the loop at lines 11–28 executes 3 times for each projection, and the worst-case complexity of the second part of the algorithm (lines 8–28) becomes  $O(3mK(K(2m - 1)))$ . Thus, the overall worst-case complexity of Alg. 9 is  $O((mK)^2)$ . ■