

A Formal Approach to Deployment of Robotic Teams in an Urban-Like Environment

Yushan Chen, Xu Chu Ding, Alin Stefanescu, and Calin Belta

Abstract. We present a computational framework for automatic synthesis of control and communication strategies for a robotic team from task specifications given as regular expressions about servicing requests in an environment. Our approach is based on two main ideas. First, we extend recent results from formal synthesis of distributed systems to check for the distributability of the task specification and to generate local specifications, while accounting for the service and communication capabilities of the robots. Second, by using a technique inspired from LTL model checking, we generate individual control and communication strategies. We illustrate the method with experimental results in our Robotic Urban-Like Environment.

1 Introduction

The goal in robot motion planning and control is to be able to specify a motion task in a rich, high level language and have the robot(s) automatically convert this specification into a set of low level primitives, such as feedback controllers and communication protocols, to accomplish the task [13, 5, 14]. In most of the existing works, the motion planning problem is simply specified as “go from A to B while avoiding obstacles” [13]. However, there are situations in which this is not enough to capture the nature of the task. Consider, for example, the miniature Robotic Urban-Like Environment (RULE) shown in Fig. 1, where a robot might be required to “Visit Road R_1 or Road R_2 without crossing Intersection I_3 , and then park in an available parking space,” while at same time obeying the traffic rules. Such a “rich” specification cannot be trivially converted to a sequence of “go from A to B ” primitives.

Yushan Chen · Xu Chu Ding · Calin Belta
Boston University, Boston, MA, US
e-mail: {yushanc, xcding, cbelta}@bu.edu


Alin Stefanescu
University of Pitesti, 110040 Pitesti, Romania
e-mail: alin.stefanescu@upit.ro

When several robots are available, the problem becomes even more interesting and challenging. Assume that several service requests occur at different locations in the city, and they need to be serviced subject to some temporal and logical constraints. Some of these requests can be serviced by one (possibly specific) robot, while others require the collaboration of two or more (possibly specific) robots. For example, assume that the task is to assemble a piece of machinery in location P_1 or P_2 from two components that can be found at P_3 and P_4 . The assembly requires the cooperation of two robots, and the collection of the components needs to be performed in parallel. Can we generate provably-correct individual control and communication strategies from such rich, global specifications? This is the problem that we address in this paper.

It has been advocated as far back as [1] and more recently in [15, 8, 24] that temporal logics, such as Linear Temporal Logic (LTL) and Computation Tree Logic (CTL) [6], can be used as “rich” specification languages in mobile robotics. All of the above works suggest that the corresponding formal verification (model checking) algorithms can be adapted for motion planning and controller synthesis from such specifications. Some related works show that such techniques can be extended to multi-agent systems through the use of parallel composition [18, 12] or reactive games [9]. However, such bottom-up approaches are expensive and can lead to state-space explosion even for relatively simple problems. As a result, one of the main challenges in the area of motion planning and control of distributed teams based on formal verification is to create provably-correct, top-down approaches in which a global, “rich” specification can be decomposed into local (individual) specifications, which can then be used to automatically synthesize robot control and communication strategies. In such a framework, the construction of the parallel composition of the individual motions is not necessary, and therefore the state-space explosion problem is avoided.

In this paper, we draw inspiration from the area of distributed formal synthesis [17] to develop such a top-down approach. We consider a team of robots that can move among the regions of a partitioned environment, and which have known capabilities of servicing a set of requests that can occur in the regions of the partition. Some of these requests can be serviced by a robot individually, while some require the cooperation of groups of robots. We present an algorithm that allows for the fully automatic synthesis of robot control and communication strategies from a task specification given as a regular expression over the set of requests. For simplicity of presentation, we model the environment as a graph and the robots as agents that

Fig. 1 Robotic Urban-Like Environment (RULE). Khepera III car-like robots move autonomously on streets while staying in their lanes, obeying traffic rules, and avoiding collisions.



can move between adjacent vertices and can communicate only when at particular vertices. This framework is quite general and can be used in conjunction with cell decomposition motion planning techniques [5]. In particular, by using feedback controllers for facet reachability in polytopes [10, 2], this scenario can be extended to robots with continuous dynamics moving in environments with polytopic partitions.

The contribution of this work is threefold. First, we develop a top-down computational framework for automatic deployment of mobile agents from global specifications given as regular expressions over environmental requests. This is a significant extension of our recent work [3] by enlarging the class of specifications for which a solution exists. Second, we provide a relaxation to the standard problem of distributed synthesis modulo synchronous products and language equivalence [17]. Specifically, we show how a satisfying distributed execution can be found when the global specification is only a traced-closed language, rather than a product language. This extends our previous work [22], in which we provided two heuristics for the case of asynchronous automata. Third, we implement and illustrate the computational framework in our Khepera-based Robotic Urban-Like Environment (Fig. 1). In this experimental setup, the robots can be automatically deployed from specifications given as regular expressions over requests occurring at regions in the city.

2 Preliminaries

Throughout this paper, we assume that the reader is familiar with automata theory [11, 20]. In this section, we merely review some concepts and introduce the notation.

For a set Σ , we use $|\Sigma|$ and 2^Σ to denote its cardinality and power set, respectively. A collection of subsets $\Delta = \{\Sigma_i \subseteq \Sigma, i \in I\}$ is called a *distribution* over Σ if $\bigcup_{i \in I} \Sigma_i = \Sigma$, where I is an index set. A word is a sequence of symbols from Σ . We denote Σ^* as the set of all finite words over Σ . A *language* is a set of words.

Definition 1. A finite state automaton (FSA) is a tuple $A = (Q, q_0, \Sigma, \rightarrow_A, F)$, where Q is the set of states, $q_0 \in Q$ is the initial state, Σ is the set (alphabet) of actions, $\rightarrow_A \subseteq Q \times \Sigma \times Q$ is the transition relation, and $F \subseteq Q$ is the set of accepting states. We also write $q \xrightarrow{\sigma}_A q'$ to denote $(q, \sigma, q') \in \rightarrow_A$.

We denote $\mathcal{L}(A)$ as the language accepted by an FSA A . The language of an FSA is called a *regular language*, which can be concisely represented by a *regular expression* (RE). Given an RE, an FSA accepting all and only the words satisfying the RE can be constructed by using an off-the-shelf tool, such as JFLAP [19].

Definition 2. The synchronous product of n FSAs $A_i = (Q_i, q_{0_i}, \Sigma_i, \rightarrow_{A_i}, F_i)$, denoted by $\parallel_{i=1}^n A_i$, is an FSA $A = (Q, q_0, \Sigma, \rightarrow_A, F)$, where $Q = Q_1 \times Q_2 \times \dots \times Q_n$, $q_0 = (q_{0_1}, q_{0_2}, \dots, q_{0_n})$, $\Sigma = \cup_{i=1}^n \Sigma_i$, and $F = F_1 \times F_2 \times \dots \times F_n$. The transition relation $\rightarrow_A \subseteq Q \times \Sigma \times Q$ is defined by $q \xrightarrow{\sigma}_A q'$ iff $\forall i \in I_\sigma : q[i] \xrightarrow{\sigma}_{A_i} q'[i]$ and $\forall i \notin I_\sigma : q[i] = q'[i]$, where $q[i]$ denotes the i th component of q and $I_\sigma = \{i \in \{1, \dots, n\} \mid \sigma \in \Sigma_i\}$.

For convenience, in the particular case when $\Sigma_1 = \Sigma_2 = \Sigma$, we use $A_1 \times A_2$ to denote $\prod_{i=1}^2 A_i$. Moreover, $\mathcal{L}(A \times B) = \mathcal{L}(A) \cap \mathcal{L}(B)$ ([20]). An FSA $\neg A$ is defined as an FSA that accepts the language $\overline{\mathcal{L}(A)}$ where $\overline{\mathcal{L}(A)} := \Sigma^* \setminus \mathcal{L}(A)$.

For a word $w \in \Sigma^*$ and a subset $S \subseteq \Sigma$, we denote by $w \upharpoonright_S$ the *projection* of w onto S , which is obtained by erasing all actions σ in w that do not belong to S . For a language $L \subseteq \Sigma^*$ and a subset $S \subseteq \Sigma$, we denote by $L \upharpoonright_S$ the projection of L onto S , which is given by $L \upharpoonright_S := \{w \upharpoonright_S \mid w \in L\}$. Starting from the observation that the projection of a regular language is a regular language, the projection of an FSA A on a subset $S \subseteq \Sigma$ is another FSA (denoted by $A \upharpoonright_S$) accepting the language $\mathcal{L}(A) \upharpoonright_S$, through ε -closure, determinization and minimization ([21]).

Definition 3. Given a distribution Δ of Σ , the product of a set of languages L_i over Σ_i is denoted by $\prod_{i \in I} L_i$ and defined as $\prod_{i \in I} L_i := \{w \in \Sigma^* \mid w \upharpoonright_{\Sigma_i} \in L_i \text{ for all } i \in I\}$. A *product language* over a distribution Δ of Σ is a language L such that $L = \prod_{i \in I} L_i$, where $L_i = L \upharpoonright_{\Sigma_i}$ for all $i \in I$.

Definition 4. Given a distribution Δ of Σ and $w, w' \in \Sigma^*$, we say that w is trace-equivalent to w' ($w \sim_{\Delta} w'$) iff $w \upharpoonright_{\Sigma_i} = w' \upharpoonright_{\Sigma_i}, \forall i \in I$. We denote by $[w]_{\Delta}$ the trace-equivalence class of $w \in \Sigma^*$. A *trace-closed language* over a distribution Δ of Σ is a language L such that for all $w \in L$, $[w]_{\Delta} \subseteq L$.

The class of trace-closed languages is closed under the operations of union, intersection and complementation. Note that a product language is trace-closed (but the converse is not true) ([16, 21, 23]).

3 Problem Formulation and Approach

Let

$$\mathcal{E} = (V, \rightarrow_{\mathcal{E}}) \tag{1}$$

be an environment graph, where V is the set of vertices and $\rightarrow_{\mathcal{E}} \subseteq V \times V$ is a relation modeling the set of edges. For example, \mathcal{E} can be the quotient graph of a partitioned environment, where V is a set of labels for the regions in the partition, and $\rightarrow_{\mathcal{E}}$ is the corresponding adjacency relation. In particular, V can be a set of labels for the roads, intersections, and parking spaces in an urban-like environment and $\rightarrow_{\mathcal{E}}$ can show how these are connected (see Fig. 2). Assume we have a team of mobile robots $\mathcal{A}_i, i \in I$, whose motions are restricted by \mathcal{E} , where I is a set of robot labels.

Let Σ be a set of service requests, or actions to be performed at the vertices of \mathcal{E} . The locations of the service requests are defined as a function $a : \Sigma \rightarrow V$ (i.e., different requests can occur at the same vertex but vertices do not share requests). There may be no request at some vertices of \mathcal{E} .

We model the capacity of the robots to service requests and cooperation among robots as a distribution Δ over Σ (i.e. $\cup_{i \in I} \Sigma_i = \Sigma$). Σ_i is the set of requests that can be serviced by the robot \mathcal{A}_i . For a given request $\sigma \in \Sigma$, we define $I_{\sigma} = \{i \in I \mid \sigma \in \Sigma_i\}$, i.e., I_{σ} is the set of labels of all the agents that can service request σ . The semantics

of this distribution is defined as follows. For an arbitrary request σ , if $|I_\sigma| = 1$ (i.e., there is only one agent that owns it), the agent can (and should) service the request by itself, independent of the other agents. If $|I_\sigma| > 1$, all the agents \mathcal{A}_i with $i \in I_\sigma$ must service the request simultaneously. An agent is said to service a request σ if it visits the vertex $a(\sigma)$. We assume that two or more robots can communicate only when they are at vertices at which shared requests occur.

We model the motion capabilities of each agent \mathcal{A}_i , $i \in I$ on the environment graph \mathcal{E} as a transition system T_i , defined as follows:

$$T_i = (V, v_{0_i}, \rightarrow_i, \Pi, \models_i), i \in I, \quad (2)$$

where $v_{0_i} \in V$ is the initial state representing the initial location of \mathcal{A}_i , $\rightarrow_i \subseteq V \times V$ is a reflexive transition relation satisfying $\rightarrow_i \subseteq \rightarrow_{\mathcal{E}} \cup_{v \in V} \{(v, v)\}$, $\Pi = \Sigma \cup \{\varepsilon\}$ is a finite set of observations, ε is the empty request, and $\models_i \subseteq V \times \Pi$ is a satisfaction relation where $(v, \varepsilon) \in \models_i, \forall v \in V$ and $(v, \sigma) \in \models_i, \sigma \in \Sigma_i$, iff $a(\sigma) = v$. A transition $(v, v') \in \rightarrow_i$ is also denoted by $v \rightarrow v'$. For an arbitrary state $v \in V$, we define $\Pi_v = \{\pi \in \Pi \mid (v, \pi) \in \models_i\} \in 2^\Pi$ as the set of all observations satisfied at v . A *trajectory* of T_i is a sequence $v(0)v(1) \dots v(n)$ with the property that $v(0) = v_{0_i}, v(i) \in V$, and $(v(i) \rightarrow v(i+1)), \forall i \geq 0$. We say a trajectory $v = v(0)v(1) \dots v(n)$ of T_i satisfies a word $w = w(0)w(1) \dots w(n)$ if $w(i) \in \Pi_{v(i)}, \forall i \geq 0$. In other words, the motion of robot \mathcal{A}_i is restricted by the transition relation \rightarrow_i , which captures motion constraints in addition to $\rightarrow_{\mathcal{E}}$. The locations of the requests in the environment are captured by relation \models_i . As it will become clear later, each vertex satisfying ε captures that a robot can pass through a vertex without servicing any request.

Definition 5. A *motion and service plan* (or MS plan for short) for robot $\mathcal{A}_i, i \in I$ is a word $ms_i \in (V \cup \Sigma_i)^*$ that satisfies the following conditions: (1) $ms_i(1) = v_{0_i}$, (2) if $ms_i(j) \in \Sigma_i$, then $ms_i(j-1) \in V$ and $ms_i(j) \in \Pi_{ms_i(j-1)}$ (i.e. $ms_i(j-1) = a(ms_i(j))$), for all $j > 1$ and (3) $ms_i \upharpoonright_V$ is a trajectory of T_i . A *motion plan* for robot $\mathcal{A}_i, i \in I$, defined as $m_i = ms_i \upharpoonright_V$, can be obtained from the MS plan by deleting all request entries $ms_i(j) \in \Sigma_i$. Similarly, a *service plan* for robot $\mathcal{A}_i, i \in I$, is defined as $s_i = ms_i \upharpoonright_{\Sigma_i}$, can be obtained from the MS plan by deleting all motion entries $ms_i(j) \in V$.

The semantics of an MS plan is as follows. A vertex entry $ms_i(j) \in V$ means that the vertex $ms_i(j)$ should be visited. A request entry $ms_i(j) \in \Sigma_i$, means that robot \mathcal{A}_i should service request $ms_i(j)$ at vertex $ms_i(j-1)$. A request entry $ms_i(j) \in \Sigma_i$, where $|I_{ms_i(j)}| > 1$, following a vertex entry $ms_i(j-1) \in V$, triggers a *wait-and-leave* protocol (i.e. synchronization across the robots that share the same request $ms_i(j)$, where $|I_{ms_i(j)}| > 1$): while at $ms_i(j-1)$, robot \mathcal{A}_i broadcasts request $ms_i(j)$ and listens for broadcasts of $ms_i(j)$ from all agents $\mathcal{A}_j, j \in I_{ms_i(j)} \setminus \{i\}$. When they are all received, the request $ms_i(j)$ is serviced and then \mathcal{A}_i moves to the next vertex.

Remark 1. Note that one robot only needs to synchronize (communicate) with other robots that share a request σ with it, where $|I_\sigma| > 1$, before servicing this shared request. The loose synchronization enables parallel executions of individual agents (i.e. the requests that are not shared by the same robot can be serviced in parallel).

Note that by the definition (conditions (2) and (3)) of an MS plan ms_i , the motion plan $m_i = ms_i \upharpoonright_V$ is a trajectory of T_i satisfying a word $w_i \in (\Sigma_i \cup \varepsilon)^*$, where its corresponding service plan $s_i = ms_i \upharpoonright_{\Sigma_i}$ is equal to $w_i \upharpoonright_{\Sigma_i}$. We say that a word s_i can be *implemented* by the robot \mathcal{A}_i if there exists a MS plan ms_i such that $ms_i \upharpoonright_{\Sigma_i} = s_i$.

Given a set of service plans $\{s_i, i \in I\}$ for the robot team, there may exist many possible sequences of requests serviced by the team due to parallel executions of individual agents (we do not assume that we know the time it takes for each agent to service requests). For a given set of MS plans $ms_i, i \in I$, we denote

$$L_{MS}^{team}(\{ms_i, i \in I\}) := \prod_{i \in I} \{s_i\}, \text{ where } s_i = ms_i \upharpoonright_{\Sigma_i}, \quad (3)$$

(see Def. (3)) as the set of all possible sequences of requests serviced by the team of robots $\mathcal{A}_i, i \in I$ while they follow their individual MS plans ms_i . For simplicity of notations, we usually denote $L_{MS}^{team}(\{ms_i, i \in I\})$ as L_{MS}^{team} when there is no ambiguity. We say that the motion of the team with MS plans $\{ms_i, i \in I\}$ satisfies a specification given as an RE ϕ over Σ if $L_{MS}^{team} \neq \emptyset$ and all words in L_{MS}^{team} satisfy ϕ (i.e. $L_{MS}^{team} \subseteq \mathcal{L}(A)$, where A is an FSA accepting only the words satisfying ϕ). We are now ready to formulate the main problem:

Problem 1. Given a team of agents $\mathcal{A}_i, i \in I$ with motion capabilities T_i (Eqn. (2)), a set of service requests Σ , a task specification ϕ in the form of an RE over Σ , and a distribution Δ over Σ , find a set of MS plans $\{ms_i, i \in I\}$ such that the motion of the team satisfies ϕ .

Remark 2. For a set of MS plans, the corresponding L_{MS}^{team} could be an empty set by the definition of product of languages (since there may not exist a word $w \in \Sigma^*$ such that $w \upharpoonright_{\Sigma_i} = s_i, \forall i \in I$). In practice, this case corresponds to a scenario where one (or more) agent waits indefinitely for other agents to service a request σ that is shared among these agents. For example, if σ does not appear in the service plan of one of the agent who owns σ but it appears in the service plans of some other agents, then all those agents will be stuck in a “deadlock” state and wait indefinitely. When such a deadlock scenario occurs, the motion of the team does not satisfy the specification.

Remark 3. We made some apparently restrictive assumptions in the formulation of Prob. 1: we assumed that the vertices do not share requests and that the robots can communicate only when they are in the same vertex. They are made for the simplicity of notation. To relax the first assumption, we can use a relation instead of a function to define the locations of requests. The second assumption can be relaxed by introducing a communication relation on V (i.e. a communication graph).

In the case that Prob. 1 has a solution, for each MS plan ms_i , a robot generates a *control and communication strategy*, which is a finite sequence of control primitives, interrupts, and communication protocols. To guarantee the uniqueness of this strategy, we assume that each robot is equipped with a set of motion primitives (feedback controllers), such that the selection of a motion primitive at a vertex uniquely determines the next vertex, given that the robot is properly initialized and the history of visited vertices is known. In other words, we assume that \mathcal{A}_i can follow any trajectory of T_i (see Sec. 5).

Our approach to solve Prob. 1 can be summarized as follows. We first generate an “implementable” FSA for each robot, which captures all the possible service plans that can be implemented by the robot (Sec. 4.1). Then, if the language satisfying the global specification ϕ is trace-closed, we generate a solution to the problem. Otherwise, we attempt to construct an FSA whose language is trace-closed and satisfies the global specification. If we succeed (the language of this FSA is not empty), then we use it to generate a solution (Sec. 4.2.) Our overall approach is summarized as a provably-correct algorithm in Sec. 4.2.

In our previous work [3], we provided a solution to Prob. 1 by following the “standard” approach to distributed synthesis modulo synchronous products and language equivalence [17]. As a result, our approach was conservative, since we could only generate a solution for the particular case when the language satisfying ϕ was a product language (Def. 3). In this paper, we show that we can find a solution to Prob. 1 if the language satisfying ϕ is trace-closed. Since trace-closed languages are less restrictive than product languages (*i.e.* product languages are trace-closed but not vice versa), we significantly reduce the conservatism from our previous approach. In addition, our current approach is less expensive. Indeed, checking whether a language is trace-closed is linear in the size of the FSA accepting the language, while checking whether a language is a product language is PSPACE-complete [21].

4 Synthesis of Local MS Plans from the Global Specification

We omit all the proofs in this section due to space limitations. They are available in our detailed technical report [4].

4.1 Synthesis of the Local Implementable Specifications

We begin with the conversion of the specification ϕ over Σ to a minimal and deterministic FSA $A = (Q, q_0, \Sigma, \rightarrow, F)$, which accepts exactly the language over Σ that satisfies ϕ (using JFLAP [19]). We call A the *global* specification. Given the distribution Δ , we assign requests to each agent. Specifically, we construct a set of projected FSAs $A_i = (Q_i, q_{0_i}, \Sigma_i, \rightarrow_{A_i}, F_i)$ whose languages are the projections of $\mathcal{L}(A)$ onto the local alphabets Σ_i , $i \in I$ (see Sec. 2). The projected FSAs are used as a starting point to find a solution to Prob. 1 because of the following proposition.

Proposition 1. *If a set of MS plans $\{ms_i, i \in I\}$ is a solution to Prob. 1, then its corresponding service plans $s_i = ms_i \upharpoonright_{\Sigma_i}$ are accepted words of A_i (*i.e.* $s_i \in \mathcal{L}(A_i), \forall i \in I$).*

However, to provide a provably correct solution for Prob. 1, it is not sufficient to simply choose an arbitrary accepted word from the projected FSAs A_i to be a service plan s_i . We need to make sure that (1) the service plan s_i can be implemented by robot \mathcal{A}_i and (2) all possible sequences of requests serviced by the team satisfy ϕ . To satisfy the first requirement, we aim to find an FSA A_i^E for each $i \in I$ such that the language of A_i^E equals all the accepted words of A_i that can be implemented

by the agent \mathcal{A}_i in the environment. We address the second requirement in the next sub-section.

To obtain A_i^E , we first construct a new FSA \widehat{A}_i from A_i by adding the action ε to Σ_i and self-transitions (q, ε, q) to each state $q \in Q_i$. For a robot, the action ε means that no request is serviced. We denote the set of all these self transitions by $\rightarrow_{\varepsilon_i}$. The FSA \widehat{A}_i , $i \in I$, can now be defined as:

$$\widehat{A}_i = (\widehat{Q}_i, \widehat{q}_0, \widehat{\Sigma}_i, \rightarrow_{\widehat{A}_i}, \widehat{F}_i), \quad (4)$$

where $\widehat{Q}_i = Q_i$, $\widehat{q}_0 = q_0$, $\widehat{\Sigma}_i = \Sigma_i \cup \{\varepsilon\}$, $\rightarrow_{\widehat{A}_i} = \rightarrow_{A_i} \cup \rightarrow_{\varepsilon_i}$, and $\widehat{F}_i = F_i$.

It is important to note that these self-transitions do not affect the semantics of A_i , since they mean that if no request is served by robot \mathcal{A}_i , then the state of the A_i remains the same. Given a word $\widehat{w} \in \mathcal{L}(\widehat{A}_i)$, we can obtain a word $w = \widehat{w} \upharpoonright_{\Sigma_i} \in \mathcal{L}(A_i)$. Note that the input ε corresponds to the observation ε in the transition system T_i and the set of inputs $\widehat{\Sigma}_i$ of \widehat{A}_i is a subset of the observations Π of T_i .

To restrict the trajectories of a TS T_i with a set of observations Π to the language accepted by an FSA with a set of actions $\widehat{\Sigma}_i \subseteq \Pi$, we define the following product automaton, which is inspired from LTL model checking [6]:

Definition 6. (Adapted from [8]) The product automaton $P_i = T_i \otimes \widehat{A}_i$ between the transition system $T_i = (V, v_0, \rightarrow_i, \Pi, \models_i)$ and the FSA $\widehat{A}_i = (\widehat{Q}_i, \widehat{q}_0, \widehat{\Sigma}_i, \rightarrow_{\widehat{A}_i}, \widehat{F}_i)$, is an FSA $P_i = (Q_{P_i}, q_{0_{P_i}}, \Sigma_{P_i}, \rightarrow_{P_i}, F_{P_i})$, where $Q_{P_i} = V \times \widehat{Q}_i$, $q_{0_{P_i}} = (v_0, \widehat{q}_0)$ is the set of initial states, $\Sigma_{P_i} = \widehat{\Sigma}_i \subseteq \Pi$ is the set of inputs and $F_{P_i} = V \times \widehat{F}_i$ is the set of accepting (final) states. The transition relation $\rightarrow_{P_i} \subseteq Q_{P_i} \times \Sigma_{P_i} \times Q_{P_i}$ is defined as $(v, q) \xrightarrow{\sigma_{P_i}}_{P_i} (v', q')$ iff $v \rightarrow_i v', q \xrightarrow{\sigma_{P_i}}_{\widehat{A}_i} q'$ and $\sigma_{P_i} \in \Pi_v$.

A transition $(v, q) \xrightarrow{\sigma}_{P_i} (v', q')$ of P_i exists iff $(v, v') \in \rightarrow_i$ and request σ occurs at vertex v , i.e. $a(\sigma) = v$. Transitions with input ε mean that a robot is moving from one vertex v to vertex v' (v may be equal to v') without servicing any request. $r_{P_i} = (v_i(0), \widehat{q}_i(0)) \dots (v_i(n), \widehat{q}_i(n))$, where $\widehat{q}_i(j) \in \widehat{Q}_i$, $v_i(j) \in V$ and $j \in \{1, \dots, n\}$ is a run accepted by the product automaton P_i , $i \in I$. We define the projection of r_{P_i} onto T_i as $\gamma_{T_i}(r_{P_i}) = v_i(0) \dots v_i(n)$. The following proposition shows that we can use a run of P_i to find a trajectory of T_i satisfying the local specification (a word of $\mathcal{L}(\widehat{A}_i)$).

Proposition 2. *Given any word $w_{\widehat{A}_i} \in \mathcal{L}(\widehat{A}_i)$, there exist at least one trajectory of T_i satisfying $w_{\widehat{A}_i}$ iff $w_{\widehat{A}_i} \in \mathcal{L}(P_i)$.*

Finally, we obtain A_i^E that captures $\mathcal{L}(P_i)$ by removing environment information stored in P_i . To achieve this, we collapse the states of P_i , by taking ε -closure, determinizing, and minimizing P_i . The interested readers are referred to [11] for more details about these procedures. Thus, given a word $w \in \mathcal{L}(A_i^E)$, there exists a word $w' \in \mathcal{L}(P_i)$ such that $w' \upharpoonright_{\Sigma_i} = w$. Using this fact, the following proposition shows that A_i^E captures the largest subset of $\mathcal{L}(A_i)$ which can be implemented by the robot \mathcal{A}_i in the environment.

Proposition 3. A word $w_i^E \in \mathcal{L}(A_i)$, $i \in I$, can be used to generate a MS plan ms_i for \mathcal{A}_i , such that $ms_i \upharpoonright_{\Sigma_i} = w_i^E$, if and only if $w_i^E \in \mathcal{L}(A_i^E)$.

4.2 Synthesis of Individual MS Plans

To solve Prob. 1, we need to guarantee that all possible sequences of request serviced by the team of robots following their MS plans satisfy the global specification. More specifically, we aim to find a set of service plans $\{s_i, i \in I\}$ such that $\|_{i \in I} \{s_i\} \subseteq \mathcal{L}(A)$ and $\|_{i \in I} \{s_i\} \neq \emptyset$. The following proposition shows that a trace-closed language is sufficient to satisfy this requirement and provide a solution to Prob. 1:

Proposition 4. Given a language L and a distribution Δ of Σ , if L is a trace-closed language and $w \in L$, then $\|_{i \in I} \{w \upharpoonright_{\Sigma_i}\} \subseteq L$.

Our approach aims to construct an FSA A_G whose language is both trace-closed and included in $\mathcal{L}(A)$. By Prop. 4, an arbitrary word accepted by A_G can be used to generate a set of service plans satisfying the desired requirement by projecting this word onto the given distribution Δ . Furthermore, we use the synchronous product (SP) of the local implementable specifications generated in the previous sub-section to ensure that the obtained service plans can be implemented by individual agents. This is achieved by taking product of automata, which produces intersection of regular languages.

Specifically, to find A_G , we first check if $\mathcal{L}(A)$ is trace-closed. An algorithm that checks this property for an arbitrary FSA can be found in [4]. If $\mathcal{L}(A)$ is trace-closed, then we define $A_G = A \times \|_{i \in I} A_i^E$. Otherwise, we define $A_G = \neg(\|_{i \in I} B_i) \times \|_{i \in I} A_i^E$, where $B_i = B \upharpoonright_{\Sigma_i}$ and $B = \|_{i \in I} A_i^E \times (\neg A)$. In this second case, A_G is constructed specifically to remove words $w \in \mathcal{L}(\|_{i \in I} A_i^E)$ that cannot be used to generate desired individual service plans for the robots (*i.e.*, $\|_{i \in I} \{s_i = w \upharpoonright_{\Sigma_i}\} \not\subseteq \mathcal{L}(A)$). The following proposition shows that A_G satisfies the desired requirement.

Proposition 5. $\mathcal{L}(A_G)$ is a trace-closed language and $\mathcal{L}(A_G) \subseteq \mathcal{L}(A)$.

If $\mathcal{L}(A_G)$ is not empty, then a solution to Prob. 1 can be found by picking any accepted word of A_G . In this paper, we obtain this word w_g by using a backward reachability search starting from the set of accepting states and ending at the initial state. In a particular application, any optimization criterion can be used. Once obtained, w_g is projected onto the given distribution to generate a set of MS plans.

The overall approach proposed in this section is summarized in Alg. 1. In the following theorem, we show that the solution obtained by Alg. 1 is provably correct.

Theorem 1. If $\mathcal{L}(A_G) \neq \emptyset$, then Alg.1 returns a solution to Prob. 1, *i.e.*, a set of MS plans $\{ms_i, i \in I\}$ such that $L_{MS}^{team} \subseteq \mathcal{L}(A)$ and $L_{MS}^{team} \neq \emptyset$.

Remark 4 (Completeness). In the case that $\mathcal{L}(A)$ is trace-closed, our approach is complete in the sense that we find a solution to Prob. 1 if one exists. This follows

Algorithm 1. Construction of a set of MS plans from a global specification

- Input:** A RE ϕ , a distribution Δ , and a set of TS $\{T_i = (V, v_0, \rightarrow_i, \Pi, \models_i), i \in I\}$
- 1: Convert ϕ to a deterministic and minimal FSA A and construct $\{A_i, i \in I\}$ ($A_i = A \upharpoonright_{\Sigma_i}, \forall i \in I$)
 - 2: Construct $\{\widehat{A}_i, i \in I\}$ from $\{A_i, i \in I\}$ (Eqn. 4) and $\{P_i = \widehat{A}_i \otimes T_i, i \in I\}$ (Def. 6)
 - 3: Take ε -closure, determinize, and minimize P_i to obtain $\{A_i^E, i \in I\}$, where $\mathcal{L}(A_i^E) = \mathcal{L}(P_i)$
 - 4: Construct $\|_{i \in I} A_i^E$, which is the synchronous product of A_i^E
 - 5: **if** $\mathcal{L}(\|_{i \in I} A_i^E) = \emptyset$, **return** no solution exists
 - 6: **if** $\mathcal{L}(A)$ is trace-closed, $A_G = A \times \|_{i \in I} A_i^E$ **else** $A_G = \neg(\|_{i \in I} (\|_{i \in I} A_i^E \times (\neg A)) \upharpoonright_{\Sigma_i}) \times \|_{i \in I} A_i^E$
 - 7: **if** $\mathcal{L}(A_G) = \emptyset$, **return** no solution found
 - 8: Find a word $w_g \in \mathcal{L}(A_G)$ and obtain a set of local words $w_i^{loc} = w_g \upharpoonright_{\Sigma_i}$
 - 9: Construct $\{\widehat{A}_i^{loc}, i \in I\}$ (Eqn. 4) from $\{A_i^{loc}, i \in I\}$, where $\mathcal{L}(A_i^{loc}) = w_i^{loc}, \forall i \in I$
 - 10: Construct $\{P_i^{loc} = \widehat{A}_i^{loc} \otimes T_i, i \in I\}$ and find the accepted runs $\{r_{P_i^{loc}}, i \in I\}$ and the corresponding accepted words $\{w_i = w_i(0) \dots w_i(n), i \in I\}$.
 - 11: Obtain $\{r_{T_i} = \gamma_{T_i}(r_{P_i^{loc}}) = v_i(0) \dots v_i(n+1), i \in I\}$ and $\{ms_i = v_i(0)w_i(0) \dots v_i(n)w_i(n) \upharpoonright_{V \cup \Sigma_i}, i \in I\}$
 - 12: **return** a set of words $\{ms_i, i \in I\}$
-

directly from Prop. 3 and the definition of product of languages. If $\mathcal{L}(A)$ is not trace-closed, a complete solution to Prob. 1 requires one to find a non-empty trace-closed subset of $\mathcal{L}(A)$ if one exists. This problem is undecidable (the proof is in [4]). Therefore, our overall approach to Prob. 1 is not complete.

Remark 5 (Complexity). Checking if a language of a DFA A is trace-closed is linear in the size of A (this can be readily seen from the algorithm checking language trace-closedness in [4]). The overall complexity of Alg. 1 also depends on the construction of $\|_{i \in I} A_i^E$ and the size of A_G . Note that the construction of $\|_{i \in I} A_i^E$ and the size of A_G are not related to the size of the transition system T_i but only with A_i , which depends on the global DFA A and the distribution Δ . This fact substantiates the statement made in the introduction that we avoid constructing the parallel composition of the individual motions (represented by T_i) to prevent state-space explosion, and therefore our method scales well with the number of agents in the team. A detailed complexity analysis can be found in our technical report [4].

5 Automatic Deployment in the RULE

In our implementation, the global specification ϕ is first converted to the minimal DFA A by using JFLAP [19]. The rest of Alg. 1 is implemented in MATLAB: (1) we take a global DFA A , a distribution Δ and a set of transition systems T_i as inputs and output a set of individual MS plans for the robotic team; (2) we use Dijkstra's algorithm (see [7]) to find a word or a run accepted by an FSA by assuming that each transition of the FSA has default cost 1; if the algorithm fails to find an accepted

run, the language of this FSA is empty; (3) we implement the standard algorithm (see [11]) for taking ϵ -closure, determining a ϵ -NFA and minimizing a DFA. The output of Alg. 1 is then mapped to control and communication strategies (described in Sec. 3) through the use of motion primitives.

In this section, we show how our solution to Prob. 1 can be used to deploy a team of robots using a rich specification to service requests occurring in a miniature city. Our Robotic Urban-Like Environment (RULE) shown in Fig. 2 is a collection of roads, intersections, and parking lots, which are connected following a simple set of rules (*e.g.*, a road connects two (not necessarily different) intersections, the parking lots can only be located on the side of (each bound of) a road). Each intersection has traffic lights that are synchronized in the usual way. A desktop computer is used to remotely control the traffic lights through XBee wireless boards. Each parking lot consists of several parking spaces, where each parking space can accommodate exactly one car, and each parking lot has enough parking spaces to accommodate all the robots at the same time. The city is easily reconfigurable through re-taping and re-placement of the wireless traffic lights in intersections.

The robots are Khepera III miniature cars. Each car can sense when entering an intersection from a road, when entering a road from an intersection, when passing in front of a parking lot, when it is correctly parked in a parking space, and when an obstacle is dangerously close. Each car can distinguish the color of a traffic light and different parking spaces in the same parking lot. Each car is programmed with motion and communication primitives allowing it to safely drive on a road, turn in an intersection, park, and communicate with other cars. All the cars can communicate through Wi-Fi with a desktop computer, which is used as an interface to the user (*i.e.*, to enter the global specification) and to perform all the computation necessary to generate the individual control and communication strategies. Once computed, these are sent to the cars, which execute the task autonomously by interacting with the environment and by communicating with each other, if necessary. We assume that the communication protocol is deadlock-free.

Modeling RULE using the framework described in Sec. 3 proceeds as follows. The set of vertices V of the environment graph \mathcal{E} is the set of labels assigned to the roads, intersections, and parking lots (see Fig. 2). The edges in $\rightarrow_{\mathcal{E}}$ show how these regions are connected. We assume that inter-robot communication is possible




Fig. 2 The topology of the city for the case study from Sec. 5 and the labels of the roads, intersections, and parking lots

only when the robots are in the same parking lot. The motion capabilities of the (identical) robots are captured by a transition system T_i (Eqn. (2)) which has 27 vertices and 42 transitions. Note that, in reality, each vertex of T_i has associated a set of motion primitives, and each transition is triggered by a Boolean combination of interrupts. For example, at vertex R_{5l} , only one motion primitive *follow_road* is available, which allows the robot to drive on the road. There is only one possible transition from R_{5l} to I_1 , which is triggered by *at_int* AND *green_light*, where *at_int* is an interrupt generated when the robot reaches the end of a road at an intersection, and *green_light* is an interrupt generated at the green color of the traffic light.

It is important to note that, by selecting a motion primitive available at a vertex, the robot can correctly execute a run of T_i , given that it is initialized on a road. Indeed, only one motion primitive (*follow_road*) is available on a road (more details about the motion primitives can be found in [4]). In other words, MS plans defined in Sec. 3 and derived as described in Sec. 4 can be immediately implemented by a robot. It is easy to see that, under some reasonable liveness assumptions about environmental events (e.g., the traffic lights will eventually turn green), such a transition system captures the motion of each robot correctly.

In the rest of this section, we present a case study. Assume that two robots (cars), labeled as C_1 and C_2 , are available for deployment in the city with the topology from Fig. 2. Assume the set of service requests is $\Sigma = \{H_1, H_2, L_1, L_2, L_3, L_4, L_5\}$, where L_i , $i = 1, 2, 3, 4, 5$ are “light” requests, which require only one robot, and therefore should be serviced in parallel, while H_i , $i = 1, 2$ are “heavy”, and require the cooperation of the two robots. Assume that C_1 can service L_1 and L_4 and C_2 can service L_2 , L_3 and L_5 , i.e., the set of requests is distributed as $\Sigma_1 = \{L_1, L_4, H_1, H_2\}$, $\Sigma_2 = \{L_2, L_3, L_5, H_1, H_2\}$ between the two agents. Assume the requests occur at the parking lots as given by the assignment function $a(L_1) = P_1$, $a(L_2) = P_2$, $a(L_3) = P_3$, $a(L_4) = P_4$, $a(L_5) = P_1$, $a(H_1) = P_4$, and $a(H_2) = P_5$. Finally, assume that the global task specification is to service L_4 and then L_5 or first service H_1 , then both L_1 and L_2 in an arbitrary order, then H_2 , and finally both L_1 and L_3 in an arbitrary order. Formally, this specification translates to the following RE over Σ : $L_4L_5 + H_1 (L_1L_2 + L_2L_1) H_2 (L_1L_3 + L_3L_1)$.

Using Alg. 1, we generate a set of FSAs A_i^E . Since RULE is fully connected, all the words accepted by A_i can be implemented. In this example, $\mathcal{L}(A)$ is neither a product language nor a trace-closed language (e.g., for $w = L_4L_5$, we have $[w]_\Delta = \{L_4L_5, L_5L_4\}$ and hence, $[w]_\Delta \not\subseteq \mathcal{L}(A)$). Therefore, the FSA A_G is obtained as described in Sec. 4.2. We choose $w_g = H_1L_1L_2H_2L_1L_3 \in \mathcal{L}(A_G)$. The corresponding service plans for C_1 and C_2 are $s_1 = H_1L_1H_2L_1$ and $s_2 = H_1L_2H_2L_3$, respectively. The FSAs generated by Alg. 1 are shown in Fig. 3. Finally, we generate the MS plans for C_1 and C_2 . By assuming that C_1 and C_2 start in R_{2l} and R_{1l} respectively, the two MS plans are

$$ms_1 : R_{2l}I_2R_{4r}I_3R_{8r}P_4H_1R_{8r}I_4R_{5l}I_1R_{6r}P_1L_1R_{6r}I_4R_{8l}P_5H_2R_{8l}I_3R_{8r}I_4R_{5l}I_1R_{6r}P_1L_1$$

$$ms_2 : R_{1l}I_1R_{3l}I_2R_{4r}I_3R_{8r}P_4H_1R_{8r}I_4R_{5l}I_1R_{3l}I_2R_{3r}P_2L_2R_{3r}I_1R_{5r}I_4R_{8l}P_5H_2R_{8l}I_3R_{8r}I_4R_{6l}P_3L_3$$




Fig. 3 The FSAs generated by applying Alg. 1

The request entry H_1 and H_2 will trigger the wait-and-leave protocol (see Sec. 3) since they are shared by both robots. The above MS plans are then mapped to control and communication strategies through the use of motion primitives and interrupts.

To demonstrate that our method scales well with respect to the number of agents in the team, we deploy 5 agents in a simulator for the RULE platform. Specifically, in this case study, the global FSA A has 9 states, the transition system T_i for each robot has 26 vertices and 41 transitions, the synchronous product $\|_{i \in I} A_i^E$ has 37 states and the determinized and minimized FSA A_G has 9 states. The generation of the MS plans for both case studies described in this section takes less than 2 seconds. The movies for the actual deployment in the RULE platform and the simulator are both available at http://hyness.bu.edu/RULE_media.html.

6 Conclusion

We presented a framework for automatic deployment of a robotic team from a specification given as a regular expression over a set of service requests occurring at known locations of a partitioned environment. Given the robot capabilities to service the requests, and the possible cooperation requirements for some requests, we find local control and communication strategies such that the global behavior of the system satisfies the given specification. We illustrate the proposed method with experimental results in our Robotic Urban-Like Environment (RULE).

We are currently pursuing several future directions. We are expanding the set of global specifications to formulas of temporal logics, such as Linear Temporal Logic, to enrich the expressiveness of the global specifications. We are also working on extensions to probabilistic frameworks. Specifically, we will use formulas of probabilistic temporal logics, such as probabilistic Linear Temporal Logic (pLTL). The satisfaction of the global specification will be guaranteed probabilistically and the deterministic transition systems will be replaced with Markov Decision Processes.

Acknowledgements. We are grateful to all reviewers for the thoughtful comments. This work was partially supported by ONR-MURI N00014-09-1051, ARO W911NF-09-1-0088, AFOSR YIP FA9550-09-1-020, and NSF CNS-0834260 at Boston University and by CNCSIS-UEFISCSU no. 7/05.08.2010 at the University of Pitesti.

References

1. Antoniotti, M., Mishra, B.: Discrete event models + temporal logic = supervisory controller: Automatic synthesis of locomotion controllers. In: IEEE International Conference on Robotics and Automation (1995)
2. Belta, C., Habets, L.: Control of a class of nonlinear systems on rectangles. *IEEE Transactions on Automatic Control* 51(11), 1749–1759 (2006)
3. Chen, Y., Birch, S., Stefanescu, A., Belta, C.: A hierarchical approach to automatic deployment of robotic teams with communication constraints. In: IEEE/RSJ International Conference on Intelligent Robots & Systems, Taipei, Taiwan, pp. 5079–5084 (2010)
4. Chen, Y., Ding, X.C., Stefanescu, A., Belta, C.: A formal approach to deployment of robotic teams in an urban-like environment. Tech. rep., Boston University (2010), hyness.bu.edu/dars
5. Choset, H., Lynch, K., Hutchinson, S., Kantor, G., Burgard, W., Kavraki, L., Thrun, S.: *Principles of Robot Motion: Theory, Algorithms, and Implementations*. MIT Press (2005)
6. Clarke, E.M., Peled, D., Grumberg, O.: *Model Checking*. MIT Press (1999)
7. Cormen, T.: *Introduction to Algorithms*. MIT press (2001)
8. Fainekos, G., Kress-Gazit, H., Pappas, G.: Hybrid controllers for path planning: A temporal logic approach. In: IEEE Conference on Decision and Control and European Control Conference, Seville, Spain, pp. 4885–4890 (2005), doi:10.1109/CDC.2005.1582935
9. Gazit, H.K., Fainekos, G., Pappas, G.J.: Where’s Waldo? Sensor-based temporal logic motion planning. In: IEEE Conference on Robotics and Automation, Rome, Italy (2007)
10. Habets, L., Collins, P., van Schuppen, J.: Reachability and control synthesis for piecewise-affine hybrid systems on simplices. *IEEE Trans. Aut. Control* 51, 938–948 (2006)
11. Hopcroft, J., Motwani, R., Ullman, J.D.: *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley (2007)
12. Kloetzer, M., Belta, C.: Automatic deployment of distributed teams of robots from temporal logic motion specifications. *IEEE Transactions on Robotics* 26(1), 48–61 (2010)
13. Latombe, J.C.: *Robot Motion Planning*. Kluwer Academic Publishers (1991)
14. LaValle, S.M.: *Planning algorithms*. Cambridge University Press, Cambridge (2006)
15. Loizou, S.G., Kyriakopoulos, K.J.: Automatic synthesis of multiagent motion tasks based on LTL specifications. In: IEEE Conference on Decision and Control, Paradise Islands, The Bahamas, pp. 153–158 (2004)
16. Mazurkiewicz, A.: Introduction to trace theory. In: *The Book of Traces*, pp. 3–41. World Scientific (1995)
17. Mukund, M.: From Global Specifications to Distributed Implementations, pp. 19–34. Kluwer Academic Publishers (2002)
18. Quottrup, M.M., Bak, T., Izadi-Zamanabadi, R.: Multi-robot motion planning: A timed automata approach. In: IEEE International Conference on Robotics and Automation, Barcelona, Spain, pp. 4417–4422 (2004)
19. Rodger, S.H., Finley, T.W.: *JFLAP: An Interactive Formal Languages and Automata Package*. Jones and Bartlett Publishers (2006)
20. Sheng, Y.: *Regular Languages*. Springer, New York (1997)
21. Stefanescu, A.: Automatic synthesis of distributed transition systems. Ph.D. thesis, University of Stuttgart (2006)

22. Ștefănescu, A., Esparza, J., Muscholl, A.: Synthesis of Distributed Algorithms Using Asynchronous Automata. In: Amadio, R.M., Lugiez, D. (eds.) CONCUR 2003. LNCS, vol. 2761, pp. 27–41. Springer, Heidelberg (2003)
23. Thiagarajan, P.S., Henriksen, J.G.: Distributed Versions of Linear Time Temporal Logic: A Trace Perspective. In: Reisig, W., Rozenberg, G. (eds.) APN 1998. LNCS, vol. 1491, pp. 643–681. Springer, Heidelberg (1998)
24. Wongpiromsarn, T., Topcu, U., Murray, R.M.: Receding horizon temporal logic planning for dynamical systems. In: IEEE Conference on Decision and Control and Chinese Control Conference, Shanghai, China, pp. 5997–6004 (2009)