

Distributed Sequential Receding Horizon Control of Multi-Agent Systems under Recurring Signal Temporal Logic

Eleftherios E. Vlahakis, Lars Lindemann, and Dimos V. Dimarogonas

Abstract—We consider the synthesis problem of a multi-agent system under signal temporal logic (STL) specifications representing bounded-time tasks that need to be satisfied recurrently over an infinite horizon. Motivated by the limited approaches to handling recurring STL systematically, we tackle the infinite-horizon control problem with a receding horizon scheme equipped with additional STL constraints that introduce minimal complexity and a backward-reachability-based terminal condition that is straightforward to construct and ensures recursive feasibility. Subsequently, we decompose the global receding horizon optimization problem into agent-level programs the objectives of which are to minimize local cost functions subject to local and joint STL constraints. We propose a scheduling policy that allows individual agents to sequentially optimize their control actions while maintaining recursive feasibility. This results in a distributed strategy that can operate online as a model predictive controller. Last, we illustrate the effectiveness of our method via a multi-agent system example assigned a surveillance task.

I. INTRODUCTION

Various applications of multi-agent systems require agents to undertake recurring tasks that must be executed over an infinite time horizon. Examples include area coverage in surveillance systems, traffic management in intelligent transportation, smart harvesting in precision agriculture, and more. Signal temporal logic (STL) [1] is an important formalism for precisely defining, among others, a wide range of such tasks. STL is equipped with quantitative semantics [2], allowing one to capture the robustness of a formula, roughly measuring the distance from the space of falsifying system runs [3]. Although formulation and verification of complex recurring multi-agent specifications are possible via STL, control synthesis from STL becomes challenging as the number of agents and the planning horizon grow.

A synthesis problem under recurring STL can be formulated as

$$\text{Minimize}_{\mathbf{u}(0), \mathbf{x}(0)} \sum_{t=0}^{\infty} \ell(x(t), u(t)) \quad (1a)$$

$$\text{subject to } x(t+1) = f(x(t), u(t)), x(0) = x_0, \quad (1b)$$

$$\mathbf{x}(0) \models \square_{[0, \infty)} \phi, \quad (1c)$$

This work was supported by the Swedish Research Council (VR), the Knut & Alice Wallenberg Foundation (KAW), the Horizon Europe Grant SymAware and the ERC Consolidator Grant LEAFHOUND.

E. E. Vlahakis and D. V. Dimarogonas are with the Division of Decision and Control Systems, School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, 10044, Stockholm, Sweden. Email: {vlahakis, dimos}@kth.se. L. Lindemann is with the Thomas Lord Department of Computer Science, Viterbi School of Engineering, University of Southern California, Los Angeles, 90089, CA, USA. Email: llindema@usc.edu.

where $\ell : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}$ is a cost function, $\mathbf{u}(0) = (u(0), u(1), \dots)$, $\mathbf{x}(0) = (x(0), x(1), \dots)$ are the optimization variables denoting admissible input and state trajectories, respectively, with $x(0)$ representing an infinite run of the system (1b) that is required to satisfy ϕ in (1c) at every time step, with \square denoting the *always* operator and ϕ an STL formula. This problem is computationally challenging and hard to solve directly due to infinite memory requirements.

Intuitively, by iteratively solving a finite-horizon variant of (1) with a sufficiently large optimization horizon and tracking past computations, a feasible input trajectory $\mathbf{u}(0)$ can be constructed by collecting the first element of the optimal solution from each iteration. Such a receding horizon approach to addressing (1) recursively has been proposed in [4] via mixed-integer programming (MIP) based on binary encoding of STL constraints. The soundness and completeness of this approach for STL formulas defined over linear predicates have resulted in approaches to reactive and robust model predictive control (MPC) of single-agent systems under STL [5]–[8]. However, due to the complexity increase caused by the use of binary variables [9], MIP-based synthesis methods are problematic when dealing with large or unbounded STL formulas and multi-agent settings.

To bypass complexity issues, the authors in [10] introduce discrete average space robustness, which allows the formulation of tractable optimization problems for control synthesis. Robustness metrics of STL that fit receding horizon schemes have also been proposed in [11], [12]. The appeal of a receding horizon approach relies on the validity that all iterations yield a solution. Although recursive feasibility is well explored for classical control objectives [13], there exist only few results in recursive feasibility of receding horizon control under STL [14], [15].

Here, we address the control synthesis problem (1) of a multi-agent system subject to a recurring STL specification, which expresses a bounded-time task that needs to be satisfied recurrently over an infinite horizon, as per (1c). We introduce a receding horizon scheme to attack (1) iteratively and propose additional constraints and terminal conditions that guarantee recursive feasibility enabling an MPC implementation. Subsequently, we decompose the receding horizon control problem into smaller agent-level problems and propose a distributed receding horizon scheme that iteratively addresses (1), preserving recursive feasibility. This is facilitated by a scheduling policy that allows agents to sequentially optimize their actions, adhering to the couplings of the global STL formula.

The remainder of the paper is organized as follows. Our

MAS setup is introduced in Sec. II. The proposed receding horizon scheme and its distributed version are in Sec. III. A numerical example is in Sec. IV, whereas concluding remarks are discussed in Sec. V.

II. PRELIMINARIES AND MULTI-AGENT SETUP

A. Notation

The sets of real numbers and nonnegative integers are \mathbb{R} and \mathbb{N} , respectively. Let $N \in \mathbb{N}$. Then, $\mathbb{N}_{[0,N]} = \{0, 1, \dots, N\}$. The transpose of ξ is ξ^\top . The identity matrix is $I_n \in \mathbb{R}^{n \times n}$. The j th element of a vector a is denoted as $(a)_j$. Let x_1, \dots, x_n be vectors not necessarily of identical dimensions. Then, $x = (x_1, \dots, x_n) = [x_1^\top \dots x_n^\top]^\top$. We denote by $\mathbf{x}(a : b) = (x(a), \dots, x(b))$ an aggregate vector consisting of $x(t)$, $t \in \mathbb{N}_{[a,b]}$, representing a trajectory. When it is clear from the context, we write $\mathbf{x}(t)$, omitting the endpoint. Let $x_i(t)$, for $t \in \mathbb{N}_{[0,N]}$ and $i \in \mathbb{N}_{[1,M]}$. Then, $\mathbf{x}(0 : N) = (x(0), \dots, x(N))$ denotes an aggregate trajectory when $x(t) = (x_1(t), \dots, x_M(t))$, $t \in \mathbb{N}_{[0,N]}$. We denote by $x(t_\tau)$ the prediction of $x(t + \tau)$ carried out at time t . We also denote by $\mathbf{x}(t_{a:b})$ the predicted trajectory $(x(t_a), \dots, x(t_b))$. The halfspace representation of a convex compact polytopic set is $\mathcal{Y} = \{y \mid G_y y \leq g_y\}$ with the inequalities applied elementwise. The cardinality of \mathcal{V} is $|\mathcal{V}|$. The Minkowski sum of $S_1 \subseteq \mathbb{R}^n$ and $S_2 \subseteq \mathbb{R}^n$ is $S_1 \oplus S_2 = \{s_1 + s_2 \mid s_1 \in S_1, s_2 \in S_2\}$. The Kronecker product of $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{p \times q}$ is $A \otimes B \in \mathbb{R}^{np \times mq}$. The remainder of the division of a by b is $\text{mod}(a, b)$.

B. Signal Temporal Logic

We consider STL formulas with syntax

$$\phi := \top \mid \pi \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 U_{[a,b]} \phi_2, \quad (2)$$

where $\pi := (\mu(x) \geq 0)$ is a predicate, $\mu : \mathbb{R}^n \rightarrow \mathbb{R}$ being a predicate function, and ϕ_1, ϕ_2 are STL formulas, which are specified recursively using predicates π , logical operators \top, \neg , and \wedge , and the *until* temporal operator U , with $[a, b] \equiv \mathbb{N}_{[a,b]}$. Using the syntax in (2), we may define \vee (*or*), \diamond (*eventually*) and \square (*always*) operators, e.g., $\phi_1 \vee \phi_2 = \neg(\neg\phi_1 \wedge \neg\phi_2)$, $\diamond_{[a,b]} \phi = \top U_{[a,b]} \phi$, and $\square_{[a,b]} \phi = \neg \diamond_{[a,b]} \neg \phi$.

We denote by $\mathbf{x}(t) \models \phi$ the satisfaction of the formula ϕ defined over the trajectory $\mathbf{x}(t) = (x(t), x(t+1), \dots)$. The validity of an STL formula ϕ defined over a trajectory $\mathbf{x}(t)$ is verified as follows: $\mathbf{x}(t) \models \pi \Leftrightarrow \mu(x(t)) \geq 0$, $\mathbf{x}(t) \models \neg\phi \Leftrightarrow \neg(\mathbf{x}(t) \models \phi)$, $\mathbf{x}(t) \models \phi_1 \wedge \phi_2 \Leftrightarrow \mathbf{x}(t) \models \phi_1 \wedge \mathbf{x}(t) \models \phi_2$, $\mathbf{x}(t) \models \phi_1 U_{[a,b]} \phi_2 \Leftrightarrow \exists \tau \in t \oplus \mathbb{N}_{[a,b]}$, s.t. $\mathbf{x}(\tau) \models \phi_2 \wedge \forall \tau' \in \mathbb{N}_{[t,\tau]}$, $\mathbf{x}(\tau') \models \phi_1$.

Based on the Boolean semantics introduced above, the horizon of a formula is recursively defined as [1]: $N^\pi = 0$, $N^{\neg\phi} = N^\phi$, $N^{\phi_1 \wedge \phi_2} = \max(N^{\phi_1}, N^{\phi_2})$, $N^{\phi_1 U_{[a,b]} \phi_2} = b + \max(N^{\phi_1}, N^{\phi_2})$.

STL is equipped with quantitative metrics for assessing the robustness of a formula [2]. A scalar-valued function $\rho^\phi : \mathbb{R}^n \times \dots \times \mathbb{R}^n \rightarrow \mathbb{R}$ indicates how robustly a trajectory $\mathbf{x}(t)$ satisfies a formula ϕ . The robustness function is defined recursively as follows: $\rho^\pi(\mathbf{x}(t)) = \mu(x(t))$,

$$\begin{aligned} \rho^{\neg\phi}(\mathbf{x}(t)) &= -\rho^\phi(\mathbf{x}(t)), \\ \rho^{\phi_1 \wedge \phi_2}(\mathbf{x}(t)) &= \min(\rho^{\phi_1}(\mathbf{x}(t)), \rho^{\phi_2}(\mathbf{x}(t))), \\ \rho^{\phi_1 U_{[a,b]} \phi_2}(\mathbf{x}(t)) &= \max_{\tau \in t \oplus \mathbb{N}_{[a,b]}} (\min(Y_1(\tau), Y_2(\tau))), \\ \text{with } Y_1(\tau) &= \rho^{\phi_1}(\mathbf{x}(\tau)), Y_2(\tau) = \min_{\tau' \in \mathbb{N}_{[t,\tau]}} \rho^{\phi_2}(\mathbf{x}(\tau')), \\ \pi &\text{ being a predicate, and } \phi, \phi_1, \text{ and } \phi_2 \text{ being STL formulas.} \end{aligned}$$

C. Multi-agent system

We consider a multi-agent system (MAS) with discrete-time dynamics written as in (1b):

$$x(t+1) = f(x(t), u(t)), \quad (3)$$

where $x(t) = (x_1(t), \dots, x_M(t)) \in \mathcal{X} \subset \mathbb{R}^n$ and $u(t) = (u_1(t), \dots, u_M(t)) \in \mathcal{U} \subset \mathbb{R}^m$ are aggregate state and input vectors, respectively, with $x_i(t) \in \mathcal{X}_i \subset \mathbb{R}^{n_i}$, $u_i(t) \in \mathcal{U}_i \subset \mathbb{R}^{m_i}$, $\mathcal{X} = \mathcal{X}_1 \times \dots \times \mathcal{X}_M$, $\mathcal{U} = \mathcal{U}_1 \times \dots \times \mathcal{U}_M$, M being the number of agents, and $x(0) = x_0$ is the initial condition. We assume that the agents are open-loop dynamically decoupled, i.e., $f(x(t), u(t)) = (f_1(x_1(t), u_1(t)), \dots, f_M(x_M(t), u_M(t)))$, with $f(x(t), u(t)) : \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$ and $f_i(x_i(t), u_i(t)) : \mathbb{R}^{n_i} \times \mathbb{R}^{m_i} \rightarrow \mathbb{R}^{n_i}$, $i \in \mathbb{N}_{[1,M]}$.

Let $\mathcal{V} = \mathbb{N}_{[1,M]}$ be the set of agents' indices. The MAS is tasked with a recurring STL specification $\psi = \square_{[0,\infty)} \phi$, that is, an STL specification, ϕ , which needs to be satisfied at every time step $t \in \mathbb{N}$. The specification ϕ is a conjunctive formula, where each conjunct is an STL task involving a subset of agents $\nu \subseteq \mathcal{V}$. We call ν a clique. Let \mathcal{K}_ϕ be the set collecting all cliques induced by ϕ . Then, the MAS is subject to the recurring STL specification:

$$\psi = \square_{[0,\infty)} \phi, \quad (4a)$$

$$\phi = \bigwedge_{\nu \in \mathcal{K}_\phi} \phi_\nu, \quad (4b)$$

where a formula ϕ_ν involves individual agents (represented by one-agent cliques), pairs of agents (represented by two-agent cliques), or groups of more than two agents (indicated by multi-agent cliques), with $|\nu| = 1$, $|\nu| = 2$, or $2 < |\nu| \leq M$, respectively. Let $\pi := (\mu(y) \geq 0)$ be a predicate in ϕ , with $y \in \mathbb{R}^{n_\nu}$. The vector $y \in \mathbb{R}^{n_\nu}$ represents either an individual state vector, $x_i \in \mathbb{R}^{n_i}$, $i \in \mathbb{N}_{[1,M]}$, or an aggregate vector, $x_\nu \in \mathbb{R}^{n_\nu}$, collecting the state vectors of agents in the clique $\nu \in \mathcal{K}_\phi$.

Remark 1 Under the formulation (4), there is no restriction on the number of subtasks that are associated with a clique $\nu \in \mathcal{K}_\phi$. In fact, ϕ_ν is allowed to be any STL formula consistent with the syntax in (2).

Assumption 1 We assume that the formulas ϕ_ν , $\nu \in \mathcal{K}_\phi$, have identical horizons which we denote as N .

Remark 2 Assumption 1 is not restrictive due to the recurring structure of ψ . To see this, let, e.g., $\mathcal{K}_\phi = \{\nu', \nu''\}$ and write $\psi = \square_{[0,\infty)} (\phi_{\nu'} \wedge \phi_{\nu''})$. Let also $N^{\phi_{\nu'}}$ and $N^{\phi_{\nu''}}$ be the horizons of $\phi_{\nu'}$ and $\phi_{\nu''}$, respectively, with $N = \max\{N^{\phi_{\nu'}}, N^{\phi_{\nu''}}\}$. Then, the satisfaction of ψ may

be enforced by satisfying $\hat{\psi} = \square_{[0,\infty)}((\square_{[0,N-N^{\phi_{\nu'}}]}\phi_{\nu'}) \wedge (\square_{[0,N-N^{\phi_{\nu''}}]}\phi_{\nu''}))$, where the horizon of $\square_{[0,N-N^{\phi_{\nu}}]}\phi_{\nu}$ is $N = N - N^{\phi_{\nu}} + N^{\phi_{\nu'}}$, with $\nu = \{\nu', \nu''\}$.

D. Decomposition of STL formula ϕ

Let $\text{cl}(i) = \{\nu \in \mathcal{K}_{\phi}, \nu \ni i\}$ be the set of cliques that contain i , and assign a local STL formula to agent i as

$$\varphi_i = \bigwedge_{q \in \text{cl}(i)} \phi_q. \quad (5)$$

Let $q \in \text{cl}(i)$, where $q = (i_1, \dots, i, \dots, i_{q'})$, with $q' + 1 = |q|$. Let a trajectory $\mathbf{x}_q(t) = (x_q(t), x_q(t+1), \dots)$, where $x_q(t) = (x_{i_1}(t), \dots, x_i(t), \dots, x_{i_{q'}}(t))$, with $t \in \mathbb{N}$, and the order $i_1 < \dots < i < \dots < i_{q'}$ being specified by the lexicographic ordering of the set $\mathcal{V} = \mathbb{N}_{[1,M]}$. Then, we denote by $\mathbf{x}_q(t) \models \phi_q$ the satisfaction of ϕ_q verified over the aggregate trajectory $\mathbf{x}_q(t)$.

Remark 3 Comparing (4b) and (5), one may easily verify that $\mathbf{x}(t) \models \bigwedge_{i \in \mathcal{V}} \varphi_i$ implies $\mathbf{x}(t) \models \phi$.

III. MAIN RESULTS

A. Receding horizon control

In this section, we propose a receding-horizon variant of (1) by incorporating additional constraints and terminal conditions that enforce trajectories to exhibit a recurring pattern. Next, we recall the following definitions.

Definition 1 For the system (3), we denote the precursor set to a set $\mathcal{S} \subset \mathcal{X}$ as follows

$$\text{Pre}(\mathcal{S}) = \{x \in \mathbb{R}^n \mid \exists u \in \mathcal{U} \text{ s.t. } f(x, u) \in \mathcal{S}\}. \quad (6)$$

Definition 2 (one-Step Controllable Set) For the admissible set \mathcal{X} and a given target set $\mathcal{S} \subseteq \mathcal{X}$ the one-step controllable set $\mathcal{C}_1(\mathcal{S})$ of the system (3) is defined as

$$\mathcal{C}_1(\mathcal{S}) = \text{Pre}(\mathcal{S}) \cap \mathcal{X}. \quad (7)$$

The set $\mathcal{C}_1(\mathcal{S})$ collects all the states of (3) that can be driven to the target set \mathcal{S} in one step. We will now introduce a receding horizon version of (1) and state its recursive feasibility property. At time $t \in \mathbb{N}$,

$$\text{Min.}_{\substack{\mathbf{u}(t_0:N) \\ \mathbf{x}(t_0:N+1)}} \sum_{i=1}^M \left(\sum_{k=0}^N \ell_i(x_i(t_k), u_i(t_k)) + V_i(x_i(t_{N+1})) \right) \quad (8a)$$

$$\text{s.t. } x(t_{k+1}) = f(x(t_k), u(t_k)), \quad k \in \mathbb{N}_{[0,N]}, \quad (8b)$$

$$(\mathbf{x}(t-N+1 : t-1), \mathbf{x}(t_0:1)) \models \phi, \quad (8c)$$

\vdots

$$(\mathbf{x}(t-1), \mathbf{x}(t_0:N-1)) \models \phi, \quad (8d)$$

$$\mathbf{x}(t_0:N) \models \phi, \quad (8e)$$

$$\mathbf{x}(t_1:N+1) \models \phi, \quad (8f)$$

\vdots

$$(\mathbf{x}(t_{N+1}), \mathbf{x}(t_1:N)) \models \phi, \quad (8g)$$

$$(\mathbf{x}(t_1:N), \mathbf{x}(t_0)) \models \phi, \quad (8h)$$

\vdots

$$(\mathbf{x}(t_N), \mathbf{x}(t_0:N-1)) \models \phi, \quad (8i)$$

$$x(t_0) = x(t), \quad (8j)$$

$$x(t_k), x(t_{N+1}) \in \mathcal{X}, \quad u(t_k) \in \mathcal{U}, \quad k \in \mathbb{N}_{[0,N]}, \quad (8k)$$

$$x(t_N) \in \mathcal{C}_1(x(t)), \quad (8l)$$

where $\ell_i(x_i(t_k), u_i(t_k))$, $V_i(x(t_{N+1}))$, with $\ell_i : \mathbb{R}^{n_i} \times \mathbb{R}^{m_i} \rightarrow \mathbb{R}_{\geq 0}$, $V_i : \mathbb{R}^{n_i} \rightarrow \mathbb{R}_{\geq 0}$, represent agent-level decomposition of the cost function in (1a), $x(t) = (x_1(t), \dots, x_M(t))$, $\phi = \bigwedge_{\nu \in \mathcal{K}_{\phi}} \phi_{\nu}$, N is the horizon of ϕ , and $\mathcal{C}_1(x(t))$ in (8l) is the one-step controllable set of system (3). The constraints in (8c)-(8d) consider trajectories composed of historical data, the current state and predicted states, whereas the one in (8e) considers a trajectory originating at $x(t)$. The constraints in (8f)-(8g) enforce that any trajectory consisting of $N+1$ successive elements of $(\mathbf{x}(t_1:N+1), \mathbf{x}(t_1:N))$ satisfies ϕ , whereas the ones in (8h)-(8i) enforce that any trajectory consisting of $N+1$ successive elements of $(\mathbf{x}(t_1:N), \mathbf{x}(t_0:N-1))$ satisfies ϕ . Last, the terminal condition in (8l) enforces recursive feasibility (see Theorem 1), and the constraints in (8k) enforce safety and input capacity constraints.

Remark 4 Given linear predicate functions, if binary encoding is employed as in [4], the constraints in (8f)-(8i) require $2N$ additional binary variables.

Theorem 1 Let the optimization problem (8) be feasible at $t = 0$. Then, it is feasible for all $t \geq 0$.

Proof: See Appendix I. \blacksquare

An immediate consequence of Theorem 1 is stated next.

Corollary 1 Let problem (8) be feasible at time $t = 0$ and denote a (not necessarily optimal) solution to it at time $t \in \mathbb{N}$ by $\mathbf{u}^*(t_0:N) = (u^*(t_0), \dots, u^*(t_N))$. Then, by applying $u(t) = u^*(t_0)$ to (3), $\mathbf{x}(t) \models \phi$ for all $t \in \mathbb{N}$, or; $\mathbf{x}(0) = (x(0), x(1), \dots)$, satisfies $\square_{[0,\infty)}\phi$, that is, $\mathbf{x}(0) \models \psi$.

Remark 5 Constructing $\mathcal{C}_1(x(t))$ is challenging for non-linear systems. As a remedy, we can replace the terminal condition in (8l) with $x(N+1|t) = x(t)$ and eliminate the redundant constraints in (8h)-(8i). This modification comes at the cost of a more restricted single-point reachability condition, but it maintains recursive feasibility.

B. Distributed receding horizon control

Next, we decompose the global control problem (8) into agent-level problems while preserving recursive feasibility. Our approach relies on the following assumption.

Assumption 2 A feasible solution to the problem (8) for $t = 0$ is known.

We are now ready to formulate the following optimization problems for all $i \in \mathbb{N}_{[1,M]}$:

$$\text{Min.}_{\mathbf{x}_i(t_{0:N+1}), \mathbf{u}_i(t_{0:N})} \sum_{k=0}^N (\ell_i(x_i(t_k), u_i(t_k))) + V_i(x_i(t_{N+1})) \quad (9a)$$

$$\text{s.t. } x_i(t_{k+1}) = f_i(x_i(t_k), u_i(t_k)), \quad k \in \mathbb{N}_{[0,N]}, \quad (9b)$$

$$(\mathbf{x}_q(t-N+1:t-1), \mathbf{x}_q(t_{0:1})) \models \phi_q, \quad (9c)$$

\vdots

$$(\mathbf{x}_q(t-1), \mathbf{x}_q(t_{0:N-1})) \models \phi_q, \quad (9d)$$

$$\mathbf{x}_q(t_{0:N}) \models \phi_q, \quad (9e)$$

$$\mathbf{x}_q(t_{1:N+1}) \models \phi_q, \quad (9f)$$

\vdots

$$(\mathbf{x}_q(t_{N+1}), \mathbf{x}_q(t_{1:N})) \models \phi_q, \quad (9g)$$

$$(\mathbf{x}_q(t_{1:N}), \mathbf{x}_q(t_0)) \models \phi_q, \quad (9h)$$

\vdots

$$(\mathbf{x}_q(t_N), \mathbf{x}_q(t_{0:N-1})) \models \phi_q, \quad (9i)$$

$$\forall q \in \text{cl}(i),$$

$$x_i(t_0) = x_i(t), \quad (9j)$$

$$x_i(t_k), x_i(t_{N+1}) \in \mathcal{X}_i, \quad u_i(t_k) \in \mathcal{U}_i, \quad k \in \mathbb{N}_{[0,N]}, \quad (9k)$$

$$x_i(t_N) \in \mathcal{C}_1(x_i(t)). \quad (9l)$$

The main idea for introducing (9) is to attack the global problem, (8), by solving smaller agent-level problems in a distributed fashion, while preserving recursive feasibility.

Lemma 1 *Let Assumption 2 hold. Then, there is always an admissible input sequence that renders (9) feasible for all $t \in \mathbb{N}$ and all agents $i \in \mathbb{N}_{[1,M]}$.*

Proof: At $t = 0$, let $\mathbf{u}(0_{0:N}) = (u(0_0), \dots, u(0_N))$, $\mathbf{x}(0_{0:N+1}) = (x(0_0), \dots, x(0_{N+1}))$ be a known feasible solution to (8), with $u(0_\tau) = (u_1(0_\tau), \dots, u_M(0_\tau))$, $\tau \in \mathbb{N}_{[0,N]}$, $x(0_{\tau'}) = (x_1(0_{\tau'}), \dots, x_M(0_{\tau'}))$, $\tau' \in \mathbb{N}_{[0,N+1]}$. By feasibility at $t = 0$, there exist inputs $\hat{u}_i(N)$ such that $f_i(x_i(0_N), \hat{u}_i(N)) = x_i(0_0)$, $i \in \mathbb{N}_{[1,M]}$. One can, then, see that there always exist $N + 1$ successive elements in the sequence $(u_i(0_0), u_i(0_1), \dots, \hat{u}_i(N), u_i(0_0), u_i(0_1), \dots, \hat{u}_i(N))$ that lie in the feasible domain of (9) for all $i \in \mathbb{N}_{[1,M]}$. ■

Lemma 1 states that the recursive feasibility property of the optimization problem (8) is preserved even if we decompose it into M agent-level problems (9). However, simultaneously solving (9) for agents belonging to the same cliques, $\nu \in \mathcal{K}_\phi$, becomes problematic, due to the difficulty to address coupling constraints in parallel. As a remedy, we propose a scheduling algorithm wherein only a subset of agents can optimize their control inputs at a given time step, whereas the remaining peers regress a feasible solution from the preceding time step. The implementation of this sequential scheduling gives rise to a distributed receding horizon control policy, while ensuring recursive feasibility.

C. Distributed sequential model predictive control

Let $\mathcal{O} = \{q \in \mathcal{V} \mid \forall q' \in \mathcal{V}, \text{cl}(q) \cap \text{cl}(q') = \emptyset\}$ be a set consisting of indices of agents that are allowed to solve their local problems (9) in parallel. Obviously, the set \mathcal{O} is not unique. We denote by $\mathcal{O}(t) \subset \mathcal{V}$ the set of agents selected at time $t \in \mathbb{N}$. We construct such a set as follows. At each time $t \in \mathbb{N}$, we initialize $\mathcal{O}(t)$ by $v(t) \in \mathcal{V}$, where $v(t) = \text{mod}(v(t-1), M) + 1$, with $M = |\mathcal{V}|$. The time dependence in $v(t)$ indicates the choice of an agent from the set \mathcal{V} at time t . For example, let $\mathcal{V} = \{1, 2, 3\}$. Then, we initialize $\mathcal{O}(0), \mathcal{O}(1), \dots$ by $v(0) = 1, v(1) = 2, v(2) = 3, v(3) = 1$, and so forth, respectively. After initialization, by going through the remaining $M - 1$ indices in \mathcal{V} , following a lexicographic order, we add an index $q \in \mathcal{V}$ to $\mathcal{O}(t)$ if $\text{cl}(q) \cap \text{cl}(q') = \emptyset$ for all $q' \in \mathcal{O}(t)$.

Procedure 1 *At time $t \in \mathbb{N}$, if $i \in \mathcal{O}(t)$, an input sequence, $\mathbf{u}_i(t_{0:N})$, is produced by solving Problem (9) (potentially optimally). If $i \notin \mathcal{O}(t)$, an input sequence, $\mathbf{u}_i(t_{0:N})$, is constructed by the preceding time step, which is feasible for Problem (9) (cf. Lemma 1).*

Based on the sequential scheduling induced by the construction of the set $\mathcal{O}(t)$, we are in a position to state the following result and its immediate consequence.

Theorem 2 *Consider Procedure 1 and let Assumption 2 hold. Then, (9) is recursively feasible for all $i \in \mathcal{V}$.*

Proof: The result relies on the construction of the set $\mathcal{O}(t)$. Let $\mathbf{u}_i(t_{0:N}), \mathbf{x}_i(t_{0:N+1})$ be feasible input and state trajectories, respectively, obtained at time $t \in \mathbb{N}$ by solving (9) for $i \in \mathcal{O}(t)$. Note that $\mathbf{x}_i(t_{0:N+1})$ does not affect the feasibility of any Problem (9) corresponding to agent $j \notin \mathcal{O}(t)$, with $\text{cl}(j) \cap \text{cl}(i) \neq \emptyset$. For such agents, a feasible input sequence for time t may be constructed by the feasible solution obtained in time step $t - 1$ (see proof of Lemma 1), which is known to all agents $i \in \mathcal{O}(t)$ at time t . Finally, recursive feasibility follows by Lemma 1. ■

Corollary 2 *Consider Procedure 1 and let Assumption 2 hold. Then, at each time $t \in \mathbb{N}$, there exist admissible local inputs $u_i(t_0)$, $i \in \mathbb{N}_{[1,M]}$ synthesizing a multi-agent input $u(t_0) = (u_1(t_0), \dots, u_M(t_0))$. These inputs, when applied to (3), yield a multi-agent state trajectory $\mathbf{x}(t)$ that satisfies ϕ for all $t \in \mathbb{N}$, that is, $\mathbf{x}(0) \models \psi$.*

A distributed MPC scheme based on the scheduling algorithm induced by the set $\mathcal{O}(t)$ and feasible solutions of problems (9), denoted as $\mathbf{u}_i^*(t_{0:N}), \mathbf{x}_i^*(t_{0:N+1})$, is implemented as follows. At time $t \in \mathbb{N}$:

- (S1) Agent i measures its local state and exchanges predictions on its receding-horizon state trajectory with agents $j \in \mathcal{V}$, where $\text{cl}(j) \cap \text{cl}(i) \neq \emptyset$, as computed in the previous time step.
- (S2) If $i \in \mathcal{O}(t)$, agent i solves (9) obtaining $\mathbf{u}_i^*(t_{0:N})$, and implements its first component, i.e., $u_i(t) = u_i^*(t_0)$. If $i \notin \mathcal{O}(t)$, $u_i(t) = u_i(t'_1)$, with $t' = t - 1$.
- (S3) At time $t + 1$, agent i repeats steps (S1) and (S2).

IV. NUMERICAL EXAMPLE

We exemplify our distributed control strategy using the MAS example considered in [10, Sec. 6], which we augment by collision avoidance among agents. The MAS consists of three identical agents, with indices, $\mathcal{V} = \{1, 2, 3\}$, with global dynamics, $x(t+1) = Ax(t) + Bu(t)$, where $x(t) = (x_1(t), x_2(t), x_3(t)) \in \mathbb{R}^{12}$, $u(t) = (u_1(t), u_2(t), u_3(t)) \in \mathbb{R}^6$, with $x_i(t) \in \mathbb{R}^4$, $u_i(t) \in \mathbb{R}^2$, $i \in \mathbb{N}_{[1,3]}$, and $A = I_6 \otimes \begin{bmatrix} 1 & 0.1 \\ 0 & 1 \end{bmatrix}$, $B = I_6 \otimes \begin{bmatrix} 0.005 \\ 1 \end{bmatrix}$. The vector $((x_i)_1, (x_i)_3) \in \mathcal{X}$ denotes the position of agent i in \mathcal{X} , where $\mathcal{X} = \{\xi \in \mathbb{R}^2 \mid [I_2 \quad -I_2]^\top \xi \leq (10, 10, 0, 0)\}$, the vector $((x_i)_2, (x_i)_4)$ collects the respective velocities, and the vector u_i collects the accelerations of agent i , where $-20 \leq (u_i)_j \leq 20$, with $j \in \mathbb{N}_{[1,2]}$, $i \in \mathbb{N}_{[1,3]}$.

The MAS is tasked with surveilling the area $\mathcal{X} \subset \mathbb{R}_{\geq 0}^2$. This specification is expressed through a recurring STL formula $\psi = \square_{[0,\infty)}\phi$, where $\phi = \phi^1 \wedge \phi^2 \wedge \phi^3$, with $\phi^1 = \phi^{11} \wedge \phi^{12} \wedge \phi^{13}$, $\phi^2 = \phi^{21} \wedge \phi^{22} \wedge \phi^{23}$, $\phi^3 = \phi^{31} \wedge \phi^{32}$. Specifically, ϕ^1 defines the admissible workspace, ϕ^2 specifies the surveillance task, and ϕ^3 specifies collision avoidance, with $\phi^{11} = \square_{[0,30]}\|((x_1)_1, (x_1)_3) - (5, 5)\|_\infty \leq 5$, $\phi^{12} = \square_{[0,30]}\|((x_2)_1, (x_2)_3) - (5, 5)\|_\infty \leq 5$, $\phi^{13} = \square_{[0,30]}\|((x_3)_1, (x_3)_3) - (5, 5)\|_\infty \leq 5$, $\phi^{21} = (\diamond_{[10,30]}\|((x_1)_1, (x_1)_3) - (5, 9)\|_\infty \leq 1) \wedge (\diamond_{[10,30]}\|((x_1)_1, (x_1)_3) - (1, 5)\|_\infty \leq 1)$, $\phi^{22} = (\diamond_{[10,30]}\|((x_2)_1, (x_2)_3) - (8, 8)\|_\infty \leq 1) \wedge (\diamond_{[10,30]}\|((x_2)_1, (x_2)_3) - (2, 2)\|_\infty \leq 1)$, $\phi^{23} = (\diamond_{[10,30]}\|((x_3)_1, (x_3)_3) - (9, 5)\|_\infty \leq 1) \wedge (\diamond_{[10,30]}\|((x_3)_1, (x_3)_3) - (5, 1)\|_\infty \leq 1)$, $\phi^{31} = \square_{[0,30]}\|((x_1)_1, (x_1)_3) - ((x_2)_1, (x_2)_3)\|_\infty \geq 0.1$, $\phi^{32} = \square_{[0,30]}\|((x_2)_1, (x_2)_3) - ((x_3)_1, (x_3)_3)\|_\infty \geq 0.1$.

To formulate local formulas, φ_i , as defined in (5), we identify the set of cliques induced by the formula ϕ as $\mathcal{K}_\phi = \{1, 2, 3, (1, 2), (2, 3)\}$, and collect the cliques associated with each agent in $\text{cl}(1) = \{1, (1, 2)\}$, $\text{cl}(2) = \{2, (1, 2), (2, 3)\}$, and $\text{cl}(3) = \{3, (2, 3)\}$. These yield local formulas $\varphi_1 = \phi^1 \wedge \phi_{12}$, $\varphi_2 = \phi^2 \wedge \phi_{12} \wedge \phi_{23}$, and $\varphi_3 = \phi^3 \wedge \phi_{23}$, where, $\phi_1 = \phi^{11} \wedge \phi^{21}$, $\phi_{12} = \phi^{31}$, $\phi_2 = \phi^{21} \wedge \phi^{22}$, $\phi_{23} = \phi^{32}$, and $\phi_3 = \phi^{13} \wedge \phi^{23}$. Next, based on the local formulas, φ_i , $i \in \mathbb{N}_{[1,3]}$, we formulate local problems (9), where we express infinity-norm-based constraints as conjunctions and disjunctions of linear inequalities and formulate the entire STL constraints in φ_i using binary encoding [4].

In our simulations, we select $\ell = \|u(t)\|_1$, $\ell_i = \|u_i(t)\|_1$, and $V_i = 0$ as global and local performance criteria, respectively. We initiate the process by solving (8) for one iteration and subsequently execute the distributed receding horizon procedure (S1)-(S3), as detailed in the previous section. Following the initial iteration, an optimization is performed by either agent 2 alone or agents 1 and 3, determined by the set $\mathcal{O}(t) \in \{\mathcal{O}_1, \mathcal{O}_2\}$, for $t \in \mathbb{N}$, where $\mathcal{O}_1 = \{2\}$ and $\mathcal{O}_2 = \{1, 3\}$. The underlying mixed-integer linear programs were formulated using the YALMIP toolbox in MATLAB [16] and solved using the GUROBI solver [17].

Starting from the chosen initial conditions, our distributed control strategy guides the MAS along a trajectory that satisfies the recurring STL formula ψ , while respecting input constraints, as shown in Fig. 2. Computational aspects for the initial centralized iteration and the subsequent distributed iterations are depicted in Fig. 1. To assess the robustness of our method, we conducted simulations by introducing velocity perturbations to the MAS for an 80-time-step duration. Despite these perturbations, the MAS effectively remains within the permissible workspace, successfully fulfilling its surveillance task, as depicted in Fig. 3. To verify both simulations, we employed the robustness function detailed in Sec. II-B, which remained nonnegative for most of the simulation period. In addition, we performed simulations in which we omitted the proposed terminal constraint. The outcomes reveal that, given the chosen input constraints, the MAS cannot successfully accomplish the surveillance task. This underscores the key role of the terminal condition in our control design.




Fig. 1: Runtime for computations on an i7-1185G7 CPU at 3.00GHz. (Left) Distributed scheme. (Right) Centralized scheme.

V. CONCLUSION

We have proposed a receding horizon control strategy to address the infinite-horizon synthesis problem of a multi-agent system under recurring STL specifications. To ensure recursive feasibility and guarantee the integrity of the proposed control scheme we have introduced additional constraints and terminal conditions that are easy to construct and intuitive for schemes under recurring tasks. By decomposing the global optimization problem into agent-level programs we implement a scheduling policy that enables individual agents to optimize their control actions sequentially. This approach yields a distributed control strategy that enables online operation, while preserving recursive feasibility. Our method hinges on the feasibility of an initial iteration carried out at the MAS level. Inspired by the proposed scheduling policy and the decomposition approach, we aim to tackle the problem in a fully distributed manner in future work.

APPENDIX I: PROOF OF THEOREM 1

We show recursive feasibility for $t > N$. The proof for $t \in \mathbb{N}_{[0,N]}$ is similar. Suppose that (8) is feasible at time $t > N$, with optimal input sequence $u^*(t_{0:N})$ and optimal state trajectory $x^*(t_{0:N+1})$. By applying $u(t) = u^*(t_0)$, the MAS is driven to $x(t+1) = f(x(t), u(t)) = x^*(t_1)$. Since $x^*(t_N) \in \mathcal{C}_1(x(t))$, there is an admissible input, $\hat{u}(t+N)$, such that $f(x^*(t_N), \hat{u}(t+N)) = x(t)$. At time $t+1$, we design admissible input sequence, $u((t+1$

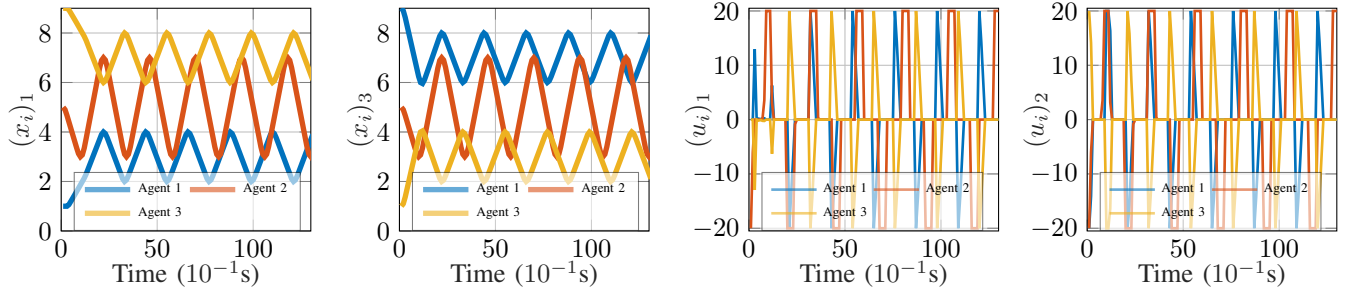


Fig. 2: Position and acceleration over a 13-second time horizon under the distributed receding horizon control scheme.

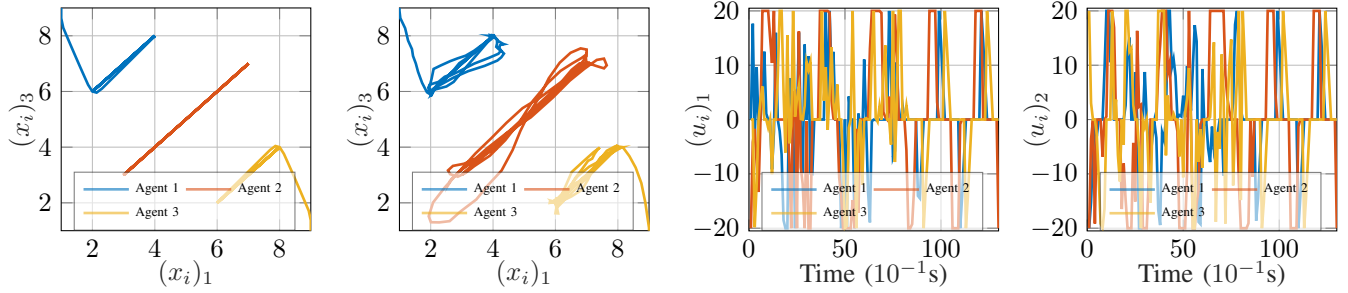


Fig. 3: (Left) Nominal and perturbed trajectories. (Right) Accelerations under random perturbations in the velocity terms.

$1)_{0:N} = (\mathbf{u}^*(t_{1:N-1}), \hat{\mathbf{u}}(t+N), \mathbf{u}^*(t_0))$, which yields $\mathbf{x}((t+1)_{0:N+1}) = (x(t+1), \mathbf{x}^*(t_{2:N}), x(t), x(t+1))$, which we check for feasibility. Constraints in (8c)-(8d), which, for time $t+1$, are written as $(x(t-N+2), x(t-N+3), \dots, \mathbf{x}^*(t_{1:2})) \models \phi$, $(x(t-N+3), x(t-N+4), \dots, \mathbf{x}^*(t_{2:3})) \models \phi$, \dots , $(x(t), \mathbf{x}^*(t_{1:N})) \models \phi$, are feasible by constraints in (8c)-(8e) corresponding to time t . Constraint in (8e) for time $t+1$, namely, $\mathbf{x}((t+1)_{0:N}) \models \phi$, which is written as $(\mathbf{x}^*(t_{1:N}), x(t)) \models \phi$, is identical to the constraint in (8h) at time t , and, hence, is feasible. Constraints in (8f)-(8g) at time $t+1$ are written as $(\mathbf{x}^*(t_{2:N+1}), \mathbf{x}^*(t_1)) \models \phi$, $(\mathbf{x}^*(t_{3:N+1}), \mathbf{x}^*(t_{1:2})) \models \phi$, \dots , $\mathbf{x}^*(t_{1:N+1}) \models \phi$, which are identical to constraints in (8f)-(8g) at time t , with a different order. Due to the choice of $x((t+1)_N) = x(t)$, constraints in (8h)-(8i) for the time step $t+1$ are identical to constraints in (8f)-(8g) at time $t+1$, and, hence, are feasible. Lastly, since $x((t+1)_{N+1}) = x(t+1)$, $x((t+1)_N) \in \mathcal{C}_1(x(t+1))$ implying that constraint in (8i) is feasible. This completes the proof.

REFERENCES

- [1] O. Maler and D. Nickovic, "Monitoring Temporal Properties of Continuous Signals," in *Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems*, Y. Lakhnech and S. Yovine, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 152–166.
- [2] A. Donzé and O. Maler, "Robust Satisfaction of Temporal Logic over Real-Valued Signals," in *8th International Conference on Formal Modeling and Analysis of Timed Systems, FORMATS, 2010*, Klosterneuburg, Austria, 2010, pp. 92–106.
- [3] G. E. Fainekos and G. J. Pappas, "Robustness of temporal logic specifications for continuous-time signals," *Theoretical Computer Science*, vol. 410, no. 42, pp. 4262–4291, 2009.
- [4] V. Raman, M. Maasoumy, A. Donze, R. M. Murray, A. Sangiovanni-Vincentelli, and S. A. Seshia, "Model predictive control with signal temporal logic specifications," in *Proceedings of the IEEE Conference on Decision and Control*. IEEE, 2014, pp. 81–87.
- [5] S. S. Farahani, V. Raman, and R. M. Murray, "Robust Model Predictive Control for Signal Temporal Logic Synthesis," *IFAC Proceedings Volumes (IFAC-PapersOnline)*, vol. 48, no. 27, pp. 323–328, 2015.
- [6] S. Sadraddini and C. Belta, "Robust temporal logic model predictive control," in *53rd Annual Allerton Conference on Communication, Control, and Computing*. IEEE, 2015, pp. 772–779.
- [7] V. Raman, A. Donzé, D. Sadigh, R. M. Murray, and S. A. Seshia, "Reactive Synthesis from Signal Temporal Logic Specifications," in *Proceedings of the 18th International Conference on Hybrid Systems: Computation and Control*, ser. HSCC '15. New York, NY, USA: ACM, 2015, pp. 239–248.
- [8] S. Ghosh, D. Sadigh, P. Nuzzo, V. Raman, A. Donzé, A. L. Sangiovanni-Vincentelli, S. S. Sastry, and S. A. Seshia, "Diagnosis and Repair for Synthesis from Signal Temporal Logic Specifications," in *Proceedings of the 19th International Conference on Hybrid Systems: Computation and Control*, ser. HSCC '16. New York, NY, USA: ACM, 2016, pp. 31–40.
- [9] V. Kurtz and H. Lin, "Mixed-Integer Programming for Signal Temporal Logic With Fewer Binary Variables," *IEEE Control Systems Letters*, vol. 6, pp. 2635–2640, 2022.
- [10] L. Lindemann and D. V. Dimarogonas, "Robust control for signal temporal logic specifications using discrete average space robustness," *Automatica*, vol. 101, pp. 377–387, 2019.
- [11] Y. V. Pant, H. Abbas, and R. Mangharam, "Smooth operator: Control using the smooth robustness of temporal logic," in *2017 IEEE Conference on Control Technol. and Applic. (CCTA)*, 2017, pp. 1235–1240.
- [12] N. Mehdipour, C.-I. Vasile, and C. Belta, "Arithmetic-geometric mean robustness for control from signal temporal logic specifications," in *2019 American Control Conference (ACC)*, 2019, pp. 1690–1695.
- [13] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, "Constrained model predictive control: Stability and optimality," *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.
- [14] S. Sadraddini and C. Belta, "Formal synthesis of control strategies for positive monotone systems," *IEEE Trans. Automat. Contr.*, vol. 64, no. 2, pp. 480–495, 2019.
- [15] M. Charitidou and D. V. Dimarogonas, "Receding horizon control with online barrier function design under signal temporal logic specifications," *IEEE Trans. Automat. Contr.*, vol. 68, no. 6, pp. 3545–3556, 2023.
- [16] J. Löfberg, "Yalmip : A toolbox for modeling and optimization in matlab," in *In Proceedings of the CACSD Conference*, Taiwan, 2004.
- [17] "Gurobi optimizer," <https://www.gurobi.com>, 2023.