

Robust MITL planning under uncertain navigation times

Alexis Linard*, Anna Gautier*, Daniel Duberg, Jana Tumova

Abstract—In environments like offices, the duration of a robot’s navigation between two locations may vary over time. For instance, reaching a kitchen may take more time during lunchtime since the corridors are crowded with people heading the same way. In this work, we address the problem of routing in such environments with tasks expressed in Metric Interval Temporal Logic (MITL) – a rich robot task specification language that allows us to capture explicit time requirements. Our objective is to find a strategy that maximizes the temporal robustness of the robot’s MITL task. As the first step towards a solution, we define a Mixed-integer linear programming approach to solving the task planning problem over a Varying Weighted Transition System, where navigation durations are deterministic but vary depending on the time of day. Then, we apply this planner to optimize for MITL temporal robustness in Markov Decision Processes, where the navigation durations between physical locations are uncertain, but the time-dependent distribution over possible delays is known. Finally, we develop a receding horizon planner for Markov Decision Processes that preserves guarantees over MITL temporal robustness. We show the scalability of our planning algorithms in simulations of robotic tasks.

Index Terms—Formal Methods, Planning Under Uncertainty, Temporal Robustness, Markov Decision Processes.

I. INTRODUCTION

We consider a scenario where a robot in an office-like environment receives various tasks over time (such as repeatedly visiting offices A, B and C three times every day, while recharging at least every 2 hours) and needs to be routed to complete them. We express tasks in Metric Interval Temporal Logic (MITL) and associate them with a priority. MITL is an extension of Linear Temporal Logic (LTL) where time intervals equip the temporal operators [1], and where the upper bound of the time intervals expresses the deadlines of the tasks. Both LTL and its timed version have been recently popular choices of robot task and motion specification language due to their rigorousness and richness [2], [3], [4], [5], [6], [7], [8]. Since MITL tasks include time constraints, the time it takes for the robot to navigate between locations will impact whether they are satisfied or not. *Temporal robustness* – a quantitative semantics originally defined for Signal Temporal Logic – measures the degree to which a strategy satisfies a specification when subjected to time shifts [9]. In other words, it indicates *how much delay a strategy can afford* while still satisfying the desired tasks.

*contributed equally. The authors are with the KTH Royal Institute of Technology, SE-100 44, Stockholm, Sweden, with the Division of Robotics, Perception and Learning. This work was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation, Digital Futures, and carried out as part of the Vinnova Competence Center for Trustworthy Edge Computing Systems and Applications at KTH Royal Institute of Technology. {linard, annagau, dduberg, tumova}@kth.se

In this paper, we consider stochastic navigation duration, where the time to move between two physical locations follows a known distribution, which, however, may change over time. For instance, if a robot has to enter the kitchen during lunchtime, it might take extra time with some positive probability. In order to successfully optimize for MITL temporal robustness under time-dependent uncertainty, as a first step, we consider the case where navigation is deterministic, but time-dependent. In this setting, a robot navigates an environment abstracted as a discrete Varying Weighted Transition System (VWTS). The states of the VWTS represent the physical locations in the environment, while the weighted transitions represent the time it takes to transit between two states. Unlike [10], these navigation times may vary depending on the time of day. Our goal is to find a *strategy*, i.e., a sequence of states in the VWTS, that satisfies MITL tasks with the least possible delays while accounting for known dynamic travel times between states.

To address the strategy synthesis for stochastically varying navigation times, we model the problem as a Markov Decision Process (MDP), where the transition function incorporates the uncertainty. The goal is to synthesize a strategy that maximizes the tasks’ expected temporal robustness. Towards this, we design a reward function that optimizes for temporal robustness. Because solving the MDP exactly can be costly, we develop a receding horizon planner for MDPs. Our receding horizon planner uses the timed VWTS model as a worst-case lookahead, thus preserving guarantees over MITL temporal robustness.

Related work includes [11] and [10], which, however, makes assumptions on static navigation duration. In [11], a vehicle’s trajectory must meet all the temporal demands within their respective deadlines. In contrast to our work, navigation times were time-independent and deterministic, and demands were limited to the syntactically co-safe fragment of LTL. [10] proposed a scalable mixed-integer encoding for a WTS with finite, deterministic, time-independent weights, subject to MITL specifications. There, the starting times and deadlines of the tasks were directly encoded in the intervals of the temporal operators. Finally, [9] defines temporal robustness for Signal Temporal Logic specifications, and requires minor adjustments to handle MITL.

Our contributions are – (A) A Mixed Integer Linear Program encoding of a VWTS with time-varying weights, thereby *integrating temporal task robustness expressed in MITL* in the optimization function. (B) A method for *optimizing MITL robustness in MDPs*, accommodating scenarios where access to a specific spatial location at a particular time may follow a probabilistic distribution rather than a

fixed time value. (C) An *evaluation on the scalability of our planning algorithms* in simulations of robotic tasks.

II. NOTATION

Let \mathbb{R} , \mathbb{Z} and \mathbb{N} be the set of real, integer and natural numbers including zero, respectively. We use a discrete notion of time throughout this paper, and time intervals are represented by $[a, b] \subset \mathbb{N}$, $a, b \in \mathbb{N}$, $a < b$.

The syntax of MITL is defined as [1]:

$$\phi := \top \mid \pi \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \mathcal{U}_{[a,b]} \phi_2$$

where $\pi \in \Pi$ is an atomic proposition; \neg and \wedge are the Boolean operators for negation and conjunction, respectively; and $\mathcal{U}_{[a,b]}$ is the temporal operator *until* over bounded interval $[a, b]$. Other Boolean operations are defined using the conjunction and negation operators to enable the full expression of propositional logic. Additional temporal operators *eventually* and *always* are defined as $\diamond_{[a,b]} \phi \equiv \top \mathcal{U}_{[a,b]} \phi$ and $\square_{[a,b]} \phi \equiv \neg \diamond_{[a,b]} \neg \phi$, respectively.

We define MITL semantics over *timed words* $\sigma = (\xi_0, t_0), (\xi_1, t_1), (\xi_2, t_2), \dots, \sigma_{t_i} = (\xi_i, t_i)$, where $\xi_i \in 2^\Pi$ is the set of atomic propositions that hold at position $i \in \mathbb{N}$ and throughout time interval $[t_i, t_{i+1}]$, $t_i \in \mathbb{N}$, $\forall i \geq 0$. We denote that σ satisfies the MITL formula ϕ with $\sigma \models \phi$, and that $\sigma_i, \sigma_{i+1}, \dots$ satisfies ϕ with $\sigma_i \models \phi$. We refer to the set of timed words as Σ . The characteristic function $\chi_\phi(\sigma) : \Sigma \rightarrow \{\pm 1\}$ indicates if a formula is satisfied [12].

$$\chi_\pi(\sigma_{t_i}) = \begin{cases} 1 & \text{if } \pi \in \xi_i \\ -1 & \text{otherwise} \end{cases}$$

$$\chi_{\neg\phi}(\sigma_{t_i}) = -\chi_\phi(\sigma_{t_i})$$

$$\chi_{\phi_1 \wedge \phi_2}(\sigma_{t_i}) = \chi_{\phi_1}(\sigma_{t_i}) \wedge \chi_{\phi_2}(\sigma_{t_i})$$

$$\chi_{\square_{[a,b]} \phi}(\sigma_{t_i}) = \bigwedge_{t' \in t_i + [a,b]} \chi_\phi(\sigma_{t'})$$

$$\chi_{\diamond_{[a,b]} \phi}(\sigma_{t_i}) = \bigvee_{t' \in t_i + [a,b]} \chi_\phi(\sigma_{t'})$$

$$\chi_{\phi_1 \mathcal{U}_{[a,b]} \phi_2}(\sigma_{t_i}) = \bigvee_{t' \in t_i + [a,b]} \left(\chi_{\phi_2}(\sigma_{t'}) \wedge \bigwedge_{t'' \in [t_i, t']} \chi_{\phi_1}(\sigma_{t''}) \right)$$

$$\chi_\phi(\sigma) = \chi_\phi(\sigma_{t_0})$$

Adapting [9], we define the (synchronous) temporal robustness $\eta : \Sigma \rightarrow \mathbb{Z}$ of an MITL formula ϕ on a word σ .

$$\eta_\phi^-(\sigma_{t_i}) = \chi_\phi(\sigma_{t_i}) \cdot \sup\{\tau \geq 0 : \forall t' \in [t_i - \tau, t_i],$$

$$\chi_\phi(\sigma_{t_i}) = \chi_\phi(\sigma_{t'})\}$$

$$\eta_\phi^+(\sigma_{t_i}) = \chi_\phi(\sigma_{t_i}) \cdot \sup\{\tau \geq 0 : \forall t' \in [t_i, t_i + \tau],$$

$$\chi_\phi(\sigma_{t_i}) = \chi_\phi(\sigma_{t'})\}$$

$$\eta_\phi^\pm(\sigma_{t_i}) = \chi_\phi(\sigma_{t_i}) \cdot \sup\{\tau \geq 0 : \forall t' \in t_i \pm [0, \tau],$$

$$\chi_\phi(\sigma_{t_i}) = \chi_\phi(\sigma_{t'})\}$$

Intuitively, temporal robustness measures how well the satisfaction of a formula ϕ holds with respect to time shifts. It quantifies the maximal amount of time that we can shift the characteristic function $\chi_\phi(\sigma_{t_i})$ of ϕ to the left ($\eta_\phi^-(\sigma_{t_i})$) or right ($\eta_\phi^+(\sigma_{t_i})$) without changing the value of $\chi_\phi(\sigma_{t_i})$. The combined notation $\eta_\phi^\pm(\sigma_{t_i})$ quantifies the maximal amount




Fig. 1. VWTS $\mathcal{A} = (S, s_0, \delta, \Pi, L, \Delta)$ shows a simplified office-like environment. States are labelled by $L = \{\text{"exit"}, \text{"off1"}, \text{"off2"}, \text{"lab"}, \text{"kitch"}\}$. Self-transitions are not shown but are assumed to have a weight of 1, and transition weights are bidirectional. Transition weights are shown along the edges, and change based on time.

of time by which we can shift $\chi_\phi(\sigma_{t_i})$ to the left and to the right, without changing the value of $\chi_\phi(\sigma_{t_i})$.

III. STRATEGY SYNTHESIS UNDER DETERMINISTIC NAVIGATION TIMES

We consider strategy synthesis for finite-state labelled Weighted Transition Systems (WTS) as a step towards handling uncertainties in navigation times. A WTS $(S, s_0, \delta, \Pi, L, C)$ is a tuple where S is a finite set of states, $s_0 \in S$ is an initial state, $\delta \subseteq S \times S$ are transition relations, Π is a finite set of atomic predicates, $L : S \rightarrow 2^\Pi$ is a labeling function associating states to atomic predicates and C is a weight function assigning weights to each of the transitions. The weight $C(s_1, s_2)$ represents the number of time steps it takes for a robot to navigate from s_1 to s_2 . Further, we define the set of adjacent states of state $s \in S$ as $Adj(s) = \{s' \in S \mid (s, s') \in \delta\}$.

For example, WTS can represent an abstraction of an office-like environment. The set of states S represent waypoints in the environment (e.g., a particular location in an office, entrance to the office, etc.); δ represents the connection between locations navigable by the robot; L is the set of labels of the office-like environment (e.g., “kitchen”, “lab”, “office”) and C the nominal number of time steps it takes for the mobile robot to navigate between waypoints.

We now generalize the classic definition of a WTS and introduce the Varying Weighted Transition Systems (VWTS). A VWTS is defined by $\mathcal{A} = (S, s_0, \delta, \Pi, L, \Delta)$, where $\Delta : (S \times \mathbb{N} \times S) \rightarrow \mathbb{N}^+$ is a *time-dependent weight function*. The weight $\Delta(s_1, t, s_2)$ represents the *number of time steps* it takes for a robot to navigate from s_1 to s_2 at timestep t . In Fig. 1, we show an example of a VWTS. In this work, we want to find a *strategy*, that is, the sequence of states of a VWTS \mathcal{A} that satisfies a given MITL specification. Further, if no satisfying solution exists, we want to find a solution that minimizes the *delay* or the *advancement* of given tasks in ϕ . In the following, we reuse some notations from [10].

A sequence of states $\mathbf{s} = s_0, s_1, \dots, s_n$ is a *path* of the VWTS \mathcal{A} if $(s_i, s_{i+1}) \in \delta \forall i \in [0, n-1]$. Each *path* \mathbf{s} is associated with a *time sequence* $\mathbb{T}(\mathbf{s}) = t_0, t_1, \dots, t_n$ where $t_0 = 0$, $t_i = t_{i-1} + \Delta(s_{i-1}, t_{i-1}, s_i)$ for $i \geq 1$, and $t_n \leq T$, where T is the planning horizon. The time t_i denotes the sum of the weights of the transitions executed, that is, the time elapsed until reaching the i -th state in the *path*

s. Further, we also associate to any path s a *timed word* $\sigma(s) = (L(s_0), t_0), (L(s_1), t_1), (L(s_2), t_2), \dots$. In other words, $\sigma(s)$ is the corresponding timed sequence of atomic propositions of path s . This *timed word* is then considered to evaluate the satisfaction of an MITL formula.

Example 1: Consider the right VWTS in Fig. 1. The path $s = s_{02}, s_{01}, s_{00}, s_{10}, s_{11}, s_{12} \dots$ yields the time sequence $\mathbb{T}(s) = 0, 3, 4, 5, 6, 12 \dots$, which results in the timed word $\sigma(s) = (\{\text{"exit"}\}, 0), (\emptyset, 3), (\{\text{"lab"}\}, 4), (\emptyset, 5), (\emptyset, 6), (\{\text{"off"}\}, 12) \dots$ where $\{\text{"exit"}\}$ is the label of s_{02} at $t_0 = 0, t_1 = 3, \dots$. Note that until the next state is reached, the label of the previous state still holds. Let us look at the specification $\square_{1,2} \text{"exit"}$. Since at times 1 and 2 the agent is transiting between states s_{02} and s_{01} , under our definition, the label of s_{02} holds until s_{01} is reached. Hence, $\square_{1,2} \text{"exit"}$ holds.

Now, consider that a robot navigating an arbitrary VWTS \mathcal{A} is given a set of tasks $D = \{(\phi_1, p_1), (\phi_2, p_2), \dots\}$, where ϕ_i is an arbitrary MITL formula and p_i its assigned priority. Priority controls the emphasis given to differing, and potentially conflicting, goals. We, therefore, define our problem as follows:

Problem 1: Given a VWTS \mathcal{A} , a set of tasks D expressed in MITL, a planning horizon $T > \|D\|$, where $\|D\|$ is the tasks' maximum horizon, we would like to find a strategy, i.e., path, that maximizes the total sum of the tasks' temporal robustness, weighted by their respective priorities, i.e.,

$$\max_s \sum_{(\phi_i, p_i) \in D} \eta_{\phi_i}(\sigma(s)) \cdot p_i \quad (1a)$$

$$\text{s.t. Transition System Constraints} \quad (1b)$$

$$\text{MITL Constraints} \quad (1c)$$

We omit specifying what specific temporal robustness (η_{ϕ}^+ , η_{ϕ}^- or η_{ϕ}^{\pm}) to use in the optimization function. Depending on the application, one might prefer optimizing one or the other. More generally, the optimization function could be any linear combination of functions.

A. Approach

1) *Encoding of the VWTS* (1b): We start by introducing a binary variable b_s^t for each state $s \in S$ and each timestep $t \in [0, T]$, where T is the planning horizon. If $b_s^t = 1$, then the optimal path visits state s at time t . The timed word resulting from the optimal path is then retrieved by $\sigma(s) = (L(s), t) \forall b_s^t$ s.t. $b_s^t = 1$. We define the constraints encoding the VWTS $\mathcal{A} = (S, s_0, \delta, \Pi, L, \Delta)$:

$$b_{s_0}^0 = 1, \quad (2a)$$

$$0 \leq \sum_{s \in S} b_s^t \leq 1 \quad \forall t \in [0, T], \quad (2b)$$

$$\sum_{s' \in \text{Adj}(s)} b_{s'}^{t+\Delta(s,t,s')} \geq b_s^t \quad \forall s \in S, \forall t \in [0, T], \quad (2c)$$

$$\sum_{s' \in \text{Adj}(s)} b_{s'}^{t+\Delta(s,t,s')} \leq 1 \quad \forall s \in S, \forall t \in [0, T]. \quad (2d)$$

Constraint (2a) establishes s_0 at $t = 0$; (2b) ensures that 0 or 1 state can be occupied at each timestep; and (2c) and (2d) enforce transition relations. While (2a) follows the encoding

presented in [10], note that (2b) differs to account for travel times that are different from 1, and that 0 or 1 state can be occupied at each timestep. Further, (2c) and (2d) differ from [10] in that they establish transition relations with dynamic time weights over time.

2) *Encoding of MITL constraints* (1c): We recursively define variables and constraints along the MITL formula's structure. We define new binary variables, z_{φ}^t , such that $z_{\varphi}^t = 1$ if and only if φ is satisfied starting from time t . Conjunction and disjunction are encoded as follows [10]:

$$z \models \bigwedge_{i=1}^n z_i \iff z \leq z_i \quad \forall i \text{ and } z \geq 1 - n + \sum_{i=1}^n z_i, \quad (3)$$

$$z \models \bigvee_{i=1}^n z_i \iff z \leq \sum_{i=1}^n z_i \text{ and } z \geq z_i \quad \forall i. \quad (4)$$

We use (3) and (4) to encode an MITL formula φ :

$$\pi \implies z_{\pi} = \sum_{s \in S | L(s) = \pi} b_s^{\tilde{q}^t}, \quad (5a)$$

$$\neg \varphi \implies z_{\neg \varphi} = 1 - z_{\varphi}, \quad (5b)$$

$$\varphi_1 \wedge \varphi_2 \implies z_{\varphi_1 \wedge \varphi_2} = z_{\varphi_1} \wedge z_{\varphi_2}, \quad (5c)$$

$$\varphi_1 \vee \varphi_2 \implies z_{\varphi_1 \vee \varphi_2} = z_{\varphi_1} \vee z_{\varphi_2}, \quad (5d)$$

$$\square_{[a,b]} \varphi \implies \bigwedge_{t' \in [t+a, t+b]} z_{\varphi}^{t'}, \quad (5e)$$

$$\diamond_{[a,b]} \varphi \implies \bigvee_{t' \in [t+a, t+b]} z_{\varphi}^{t'}, \quad (5f)$$

$$\varphi_1 \mathcal{U}_{[a,b]} \varphi_2 \implies \bigvee_{t' \in [t+a, t+b]} \left(z_{\varphi_2}^{t'} \wedge \bigwedge_{t'' \in [t, t'-1]} z_{\varphi_1}^{t''} \right). \quad (5g)$$

where \tilde{q}^t is an integer variable referring to the last time step where $b_s^{t'} = 1$ holds :

$$q^t = (q^{t-1} + 1) \cdot (1 - \sum_{s \in S} b_s^t) \quad \forall t \in [0, T] \quad (6a)$$

$$\tilde{q}^t = t - q^t \quad (6b)$$

where \tilde{q}^t is a counter variable counting for how many time steps $b_s^{t'} = 0$ holds from time t . Note that the definition of (5a) and (6a) enable the evaluation of atomic predicates when no observation can be made while transiting between 2 states, as remarked in Example 1.

3) *Encoding of MITL temporal robustness* (1a): We follow the encoding of constraints for the left and right temporal robustnesses η_{ϕ}^+ and η_{ϕ}^- presented in [9]. A *positive* value of η_{ϕ}^+ stands for how many time units $\sigma(s)$ can be advanced (or shifted to the left) while maintaining the satisfaction of ϕ . Conversely, a *negative* value of η_{ϕ}^+ stands for how many time units $\sigma(s)$ should be advanced (or shifted to the left) to obtain the satisfaction of ϕ , in other words, a *negative* value of η_{ϕ}^+ stands for the delay in which ϕ is satisfied. Concerning the right temporal robustness, a *positive* value of η_{ϕ}^- stands for how many time units $\sigma(s)$ can be postponed (or shifted to the right) while maintaining the satisfaction of ϕ . In other words, a *positive* value of η_{ϕ}^- stands for how much time one can afford to postpone completion of ϕ .




Fig. 2. MDP $\mathcal{M} = (S, A, P, T, \Pi, L)$ is simplified example with 3 states. Probabilistic transitions are shown by the tables next to each edge. Here, transitions are time-independent, but our model allows for time-dependent transitions as well.

Since the encoding of the left temporal robustness has been explicitly defined in [9], we solely define the encoding of the right temporal robustness in the following, and refer the reader to [9] or our implementation for the left temporal robustness. If $z_\varphi^t = 1$, we count the maximum number of sequential time points $t' < t$ in the past for which $z_\varphi^{t'} = 1$. If $z_\varphi^t = 0$, we then want to count the maximum number of sequential time points $t' < t$ in the past for which $z_\varphi^{t'} = 0$, and then multiply this number with -1 . To encode the right temporal robustness, we adapt the encoding proposed in [9] for the left temporal robustness, and introduce counter variables $c_{1,\varphi}^t, c_{0,\varphi}^t, \tilde{c}_{1,\varphi}^t$ and $\tilde{c}_{0,\varphi}^t$, for all $t \in [-T', 0]$, where T' is a user-defined constant standing for the maximum right temporal robustness to be calculated. Note that the encoding of the right temporal robustness also requires defining fictive states b_s^t in negative time steps ranging from $[-T', 0]$, with:

$$b_s^t = 0 \quad \forall s \in S, \forall t \in [-T', 0], \quad (7a)$$

$$z_{\phi_i}^t \quad \forall (\phi_i, p_i) \in D, \forall t \in [0, T]. \quad (7b)$$

Constraints for the right temporal robustness are defined:

$$c_{1,\varphi}^t \iff (c_{1,\varphi}^{t-1} + 1) \cdot z_\varphi^t, \quad c_{1,\varphi}^{t-T'-1} = 0, \quad (8a)$$

$$c_{0,\varphi}^t \iff (c_{0,\varphi}^{t-1} - 1) \cdot (1 - z_\varphi^t), \quad c_{0,\varphi}^{t-T'-1} = 0, \quad (8b)$$

$$\tilde{c}_{1,\varphi}^t \iff c_{1,\varphi}^t - z_\varphi^t, \quad (8c)$$

$$\tilde{c}_{0,\varphi}^t \iff c_{0,\varphi}^t + (1 - z_\varphi^t), \quad (8d)$$

$$\eta_\varphi^-(\sigma(s)) \iff \tilde{c}_{1,\varphi}^0 + \tilde{c}_{0,\varphi}^0. \quad (8e)$$

IV. STRATEGY SYNTHESIS UNDER UNCERTAIN NAVIGATION TIMES

Consider now that disturbances both vary over time and follow some (known) stochastic distribution. For instance, in an office-like environment, access to the kitchen between noon and 12:30 p.m. might take 10 minutes half of the time, and 20 minutes the other half of the time. To represent this type of uncertainty, we now consider strategy synthesis over finite-state, labeled MDPs [13]. We define a labeled MDP by $\mathcal{M} = (S, A, P, T, \Pi, L)$. As in the VWTS, S is a finite set of states, and s_0 is an initial state. Then, A is the set of actions. In this setting, actions describe either navigation from one location to another or waiting. Thus, there is one action per edge in the VWTS, all of which are represented




Fig. 3. Reachable states of the MDP from Fig. 2 after three transitions. by $A = \{a_{s_1, s_2} \mid \forall (s_1, s_2) \in \delta\}$. The transition function P encodes the stochastic delays. Formally, we define $P : S \times [0, T] \times A \times S \times [0, T] \rightarrow [0, 1]$ where $P(s, t, a, s', t')$ is the probability that, if we execute action a at time t in state s , we will arrive in state s' at time t' . Finally, as in the VWTS, T is the finite time horizon, Π is the set of atomic predicates, and L is the labeling function. Fig. 2 displays an example of a three-state MDP. A *stochastic strategy* in the MDP is defined by $\mu : S \times [0, T] \times A \rightarrow [0, 1]$. Then, $\mu(s, t, a) = p$ means that when the robot is at state s at time t , it should take action a with probability p .

A. Problem Definition and Approach

Problem 2: Given an MDP \mathcal{M} and a set of tasks D expressed in MITL, we would like to find a *strategy* μ that maximizes the *expected* total sum of the tasks' temporal robustnesses, weighted by their respective priorities, i.e.,

$$\max_{\mu} E_{\mu} \left[\sum_{(\phi_i, p_i) \in D} \eta_{\phi_i}(\sigma(S_{\mu})) \cdot p_i \right]. \quad (9a)$$

In order to succinctly represent Problem 2 with Mixed-integer linear programming, we need to define an extended MDP $\tilde{\mathcal{M}} = (\tilde{S}, A, \tilde{P}, T, \Pi, \tilde{L})$, which is augmented to include *trace history* inside the state space. Formally, $\tilde{S} = \bigcup_{t \in [0, T]} (S \times [0, T])^t$, and denotes not only the current state, but all previous states the robot has visited. For example, $(s_1, t_1), (s_2, t_2), (s_3, t_3) \in \bigcup_{t \in [0, T]} (S \times [0, T])^t$ means that the robot visited state s_1 at time t_1 , s_2 at time t_2 , and is currently at s_3 at time t_3 . For the initial state $s_0 \in S$, $(s_0, 0) \in \tilde{S}$ is the initial state in $\tilde{\mathcal{M}}$. Then, $\tilde{P} : \tilde{S} \times A \rightarrow [0, 1]$ is defined with respect to P . For $\tilde{s} = (s_1, t_1), (s_2, t_2), \dots, (s_n, t_n)$, if $P(s_n, t_n, A, s', t') = p$, then $\tilde{P}(\tilde{s}, a, \tilde{s}') = p$ for $\tilde{s}' = (s_1, t_1), (s_2, t_2), \dots, (s_n, t_n), (s', t')$. Finally, $\tilde{L} : \tilde{S} \rightarrow 2^{\Pi}$ is defined by $\tilde{L}(\tilde{s}) = L(s_n)$, for arbitrary $\tilde{s} = (s_1, t_1), (s_2, t_2), \dots, (s_n, t_n)$. Fig. 3 displays the reachable states of an MDP with trace history.

Now, we define a 2-stage Mixed-integer linear programming problem to synthesize optimal strategies with respect to (9).

1) *A reward function for MITL temporal robustness:* First, we define the reward function $\tilde{R} : \tilde{S} \rightarrow \mathbb{R}$ of our MDP $\tilde{\mathcal{M}}$. Consider an arbitrary state $\tilde{s} = (s_1, t_1), (s_2, t_2), \dots, (s_m, t_m) \in \tilde{S}$. If $t_m = T$, then $\tilde{R}(\tilde{s}) = \sum_{(\phi_i, p_i) \in D} \eta_{\phi_i}(\sigma(\tilde{s})) \cdot p_i$, where σ is defined as in Section III. This reward is explicitly calculated for each possible \tilde{s} such that $t_m = T$ through a MILP, as in [9]. For any \tilde{s} such that

$t_n \neq T$, $\tilde{R}(\tilde{s}) = 0$. This reward structure is designed so that a reward is accumulated only when an agent completes their full execution over the time horizon T , and thus, their full trace is available to compute the temporal robustness.

2) *Solving the MDP*: Next, we synthesize an optimal strategy μ by solving a Linear Program (LP) as in [13]. In practice, other MDP solvers can be used at this stage. [13] introduces a continuous variable $0 \leq o_{\tilde{s},a}^h \leq 1$ for each state $\tilde{s} \in \tilde{\mathcal{S}}$, for each action $a \in A$, and for each planning-step $h \in [0, T]$, where h represents the current number of transitions. Note that this is distinct from the current timestep, which is kept track of in the state \tilde{s} . Each variable $o_{\tilde{s},a}^h$ stands for the occupancy measure of the history-dependent MDP state: it represents the probability that, for a given synthesized strategy, the robot occupies state \tilde{s} at planning step h and takes action a . Now, the LP is:

$$\max_{\mu} \sum_{a \in A} \sum_{\tilde{s} \in \tilde{\mathcal{S}}} o_{\tilde{s},a}^T \tilde{R}(\tilde{s}) \quad (10a)$$

$$\text{s.t.} \sum_{a \in A} o_{(s_0,0),a}^0 = 1 \quad (10b)$$

$$\sum_{a \in A} o_{\tilde{s}',a'}^{h+1} = \sum_{\tilde{s} \in \tilde{\mathcal{S}}} \sum_{a \in A} o_{\tilde{s},a}^h \tilde{P}(\tilde{s}, a, \tilde{s}') \quad (10c)$$

$$\forall \tilde{s}' \in \tilde{\mathcal{S}}, \forall h \in [0, T-1]$$

Constraint (10b) ensures that the strategy occupies the initial state at $h = 0$ with probability 1. Then, (10c) encodes the transition probabilities. Once the LP is solved, the optimal strategy can be constructed from the occupancy measures.

B. Planning with a receding horizon

When planning over history-dependent MDPs, the state space of \tilde{M} is exponential in the time horizon T . As a result, it may be desirable to plan for a receding horizon $T_r < T$, and re-plan throughout execution. We now present a method of receding horizon planning that preserves guarantees over temporal robustness via a worst-case lookahead. We describe this process for right temporal robustness, but it can be modified to left temporal robustness.

The receding horizon method we present here exclusively modifies the reward function of our MDP, which is calculated explicitly before solving the MDP. Before describing this modified reward function, we first need to describe how to build a worst-case VWTS (\tilde{W}) from any MDP. \tilde{W} describes the maximum possible delays that could occur if the uncertainty in the MDP was realized in an adversarial manner. For example, if at time t , transitioning from s_1 to s_2 takes time 1 with probability 0.5 and takes time 2 with probability 0.5, the worst-case VWTS will take time 2 to transition from s_1 to s_2 at time t . Therefore, for a given MDP $\mathcal{M} = (S, A, P, T, \Pi, L)$, we build a worst-case VWTS $\tilde{W} = (S, s_0, \delta, \Pi, L, \Delta)$ where S, s_0, Π, L remain the same and δ, Δ are defined as: $(s_1, s_2) \in \delta$ if and only if there exists some $h_1, h_2 \in [0, T], a \in A$ such that $P(s_1, h_1, a, s_2, h_2) > 0$. Then for every $(s_1, s_2) \in \delta, h_1 \in [0, T]$, we define:

$$\begin{aligned} \Delta(s_1, h_1, s_2) \\ = \max\{h_2 | h_2 \in [0, T], a \in A, P(s_1, h_1, a, s_2, h_2) > 0\}. \end{aligned}$$

We can now redefine the reward function used in (10a). For a given history $\tilde{s} \in \tilde{\mathcal{S}}$ for horizon T_r , we can extend \tilde{s} to include the optimal trace in the worst-case VWTS. To do this, we define a new WTS, $\tilde{W}_{\tilde{s}} = (S, s_0, \delta, \Pi, L, \Delta_{\tilde{s}})$, where S, s_0, δ, Π, L are the same as in \tilde{W} , but the weight function Δ is modified to ensure that the only feasible trace through the transition system before time T_r is \tilde{s} . Formally:

- 1) if $h < T_r, s_1, h, s_2 \in \tilde{s}$, then $\Delta_{\tilde{s}}(s_1, h, s_2)$ is equal to the realized time delay between s_1 and s_2 in trace \tilde{s} ,
- 2) if $h < T_r, s_1, h, s_2 \notin \tilde{s}$, then $\Delta_{\tilde{s}}(s_1, h, s_2) = \infty$, and
- 3) if $h > T_r, \Delta_{\tilde{s}}(s_1, h, s_2) = \Delta_{\tilde{s}}(s_1, h, s_2)$.

We now use $\tilde{W}_{\tilde{s}}$ to define a new, worst-case lookahead reward function by $\tilde{R} : \tilde{\mathcal{S}} \rightarrow \mathbb{R}$, where $\tilde{R}(\tilde{s})$ is equal to the temporal robustness of $\tilde{W}_{\tilde{s}}$ as calculated by Section III.

Lemma 4.1: Solving the LP with $\tilde{R}(\tilde{s})$ returns the worst-case temporal robustness that could be achieved after any number of future, receding horizon replannings.

Proof: (Sketch) For notational convenience, we show this for one replanning, but the results extend to multiple replannings. The final strategy can be described by $\mu = \mu_{[0, T_r]} + \mu_{[T_r, T]}$ where $\mu_{[0, T_r]}$ describes how to plan until T_r and $\mu_{[T_r, T]}$ describes how to plan after. Strategy μ will have expected temporal robustness described by:

$$\sum_{a \in A} \sum_{\tilde{s} \in \tilde{\mathcal{S}}} o_{\tilde{s},a}^T \tilde{R}(\tilde{s}). \quad (11)$$

We can partition the T -length elements of $\tilde{\mathcal{S}}$ into sets that share a prefix of length T_r . Let $\tilde{\mathcal{S}}_T, \tilde{\mathcal{S}}_{T_r}$ be the set of T -length and T_r -length prefixes, respectively. Then we can equivalently define expected temporal robustness for μ as:

$$\sum_{a \in A} \sum_{\tilde{s}_{T_r} \in \tilde{\mathcal{S}}_{T_r}} \sum_{\tilde{s}_T \in \tilde{\mathcal{S}}_T \ni \tilde{s}_{T_r} \in \tilde{s}_T} o_{\tilde{s}_T, a}^T \tilde{R}(\tilde{s}_T). \quad (12)$$

But, for any $\tilde{s}_T \in \tilde{\mathcal{S}}_T \ni \tilde{s}_{T_r}, \tilde{R}(\tilde{s}_T) \geq \tilde{R}(\tilde{s}_{T_r})$, because the prefix for \tilde{s}_T is \tilde{s}_{T_r} , and \tilde{R} is a worst-case temporal robustness for \tilde{s}_{T_r} . Thus, the expected temporal robustness for strategy μ is at least:

$$\sum_{a \in A} \sum_{\tilde{s}_{T_r} \in \tilde{\mathcal{S}}_{T_r}} \sum_{\tilde{s}_T \in \tilde{\mathcal{S}}_T \ni \tilde{s}_{T_r} \in \tilde{s}_T} o_{\tilde{s}_T, a}^T \tilde{R}(\tilde{s}_{T_r}) \quad (13)$$

$$= \sum_{a \in A} \sum_{\tilde{s}_{T_r} \in \tilde{\mathcal{S}}_{T_r}} \tilde{R}(\tilde{s}_{T_r}) \sum_{\tilde{s}_T \in \tilde{\mathcal{S}}_T \ni \tilde{s}_{T_r} \in \tilde{s}_T} o_{\tilde{s}_T, a}^T \quad (14)$$

Finally, we note that $\sum_{\tilde{s}_T \in \tilde{\mathcal{S}}_T \ni \tilde{s}_{T_r} \in \tilde{s}_T} o_{\tilde{s}_T, a}^T$ is exactly equal to $o_{\tilde{s}_{T_r}, a}^{T_r}$, because $\{\tilde{s}_T \in \tilde{\mathcal{S}}_T \ni \tilde{s}_{T_r} \in \tilde{s}_T\}$ contains all possible extension of \tilde{s}_{T_r} . Finally, we conclude that the temporal robustness for the final strategy μ is at least as much as the temporal robustness determined before pre-planning. ■

V. EXPERIMENTS

We implemented and tested our methods in Python 3.8, using Gurobi [14] for the LP solving¹. We ran simulations on an Intel i7-8665U CPU and 32GB RAM, with a timeout of 30000 seconds.

¹https://github.com/KTH-RPL-Planiacs/mitl_task_solver_temporal_robustness




Fig. 4. Transition system of an office-like environment. Labels and self-transitions are omitted for readability. Transitions are bi-directional. For more details, see our repository¹.

TABLE I

NO. OF TASKS ($|D|$); NO. OF STATES IN THE VWTS ($|S|$); TIME HORIZON (T); ENCODING TIME OF THE LP CONSTRAINTS (t_{ENCODING}); SOLVING TIME BY THE SOLVER (t_{SOLVING}); NO. OF LP VARIABLES ($\#\text{LPVARS}$); AND NO. OF LP CONSTRAINTS ($\#\text{LPCONST}$).

$ S $	$ D $	T	t_{encoding} (s)	t_{solving} (s)	$\#\text{lpvars}$	$\#\text{lpconst}$
46	5	50	0.138	0.408	4037	7543
		100	0.268	6.353	6943	14559
		500	1.181	29.79	25851	61903
		1000	2.049	44.03	49146	104683
	10	50	0.204	0.267	4845	9929
		100	0.354	40.08	7803	18023
		500	1.453	50.00	27736	73563
		1000	2.396	298.6	51381	118643
	20	50	0.579	0.619	7076	27427
		100	0.973	293.2	10876	47383
		500	4.153	6123	33068	204683
		1000	4.870	13091	56937	233051
92	5	50	0.254	0.404	7274	12981
		100	0.479	18.23	12198	23581
		500	2.127	33.88	49723	102241
		1000	4.129	37.82	96063	195121
	10	50	0.414	0.765	8082	15367
		100	0.588	14.04	12993	26951
		500	2.246	23.63	50188	104941
		1000	4.525	91.30	98358	208229
	20	50	0.827	1.629	10313	32865
		100	1.314	23723	16492	52193
		500	4.019	>30000	56361	167121
		1000	6.037	>30000	102356	259121

A. Deterministic Navigation Times

We ran experiments varying the number of states, the number of tasks, and the horizon on a VWTS with transitions from the office-like environment described in Fig. 4. Table I shows that, as the time horizon T and the number of tasks D increase, the time it takes to solve the MILP and yield a strategy increases. Our results demonstrate the scalability of the *deterministic navigation times* scenario. One could achieve day-long task planning even with time steps of 1 minute. One can easily adjust the time resolution and the planning time window.

B. Uncertain Navigation Times

We ran experiments varying the number of states, the number of tasks, the horizon, and the receding horizon on a uncertain navigation time MDP with transitions from the office-like environment described in Fig. 4. The results in Table II show that the time to encode the MDP is significantly more cumbersome than the time to solve the

TABLE II

NO. OF TASKS ($|D|$); NO. OF STATES IN THE MDP ($|S|$); TIME HORIZON (T); RECEDING HORIZON (T_r); ENCODING TIME OF LP CONSTRAINTS (t_{ENCODING}); SOLVING TIME BY SOLVER (t_{SOLVING}); NO. OF LP VARIABLES ($\#\text{LPVARS}$); AND NO. OF LP CONSTRAINTS ($\#\text{LPCONST}$).

$ S $	$ D $	T	T_r	t_{encoding} (s)	t_{solving} (s)	$\#\text{lpvars}$	$\#\text{lpconst}$
46	2	25	5	28.86	0.054	1186	161
			7	479.6	0.364	14596	1964
		50	5	255.7	0.055	1186	161
			7	2271	0.317	14596	1964
	5	25	5	36.08	0.037	1186	161
			7	578.0	0.379	14596	1964
		50	5	70.66	0.085	1186	161
			7	906.7	0.342	14596	1964

MDP. The encoding time includes explicitly defining the reward function, which relies on solving a series of calls to the VWTS MITL temporal robustness MILP. As the receding horizon T_r increases, the number of calls to that MILP increases exponentially. As the time horizon T increases and the number of tasks D increases, the time it takes to execute that MILP increases.

VI. CONCLUSIONS AND FUTURE WORK

We developed a planning methodology for optimizing MITL temporal robustness in scenarios where robot navigation times are uncertain. For real-world applications, the algorithm considering the MDP model finds its most beneficial use within relatively short planning horizons. This approach is particularly suited for situations where frequent replanning with fewer task demands is necessary. While our study primarily focused on a single robot navigating an office-like environment with multiple tasks and uncertain navigation times, we believe our method has broader applicability. In future, we aim to extend it to scenarios involving robots operating in dynamic environments, such as navigation with varying levels of crowds, or subject to travel delays induced by human activity [15]. We aim to learn a spatio-temporal model of human occupation from temporal sequences [16]. We would use such an activity distribution model, as well as people's typical schedules for daily activities, as input to our planner. In these cases, the stochastic nature of navigation durations becomes crucial, as the robot must adapt to unpredictable delays during task execution. Additionally, we aim to extend our strategy synthesis to scenarios involving multiple robots. Finally, we believe that temporal robustness is a good measure of a strategy's resilience in accommodating time shifts while still satisfying temporal task requirements. It provides a quantitative measure of adaptability, particularly important in addressing navigational challenges where uncertain navigation times can affect a robot's ability to meet task deadlines and efficiently prioritize tasks. Furthermore, the optimization function can be customized to linear functions, depending on specific optimization objectives and applications.

ACKNOWLEDGEMENT

We would like to thank Patric Jensfelt for his support and valuable discussions.

REFERENCES

- [1] R. Koymans, "Specifying real-time properties with metric temporal logic," *Real-time systems*, vol. 2, no. 4, pp. 255–299, 1990.
- [2] G. Fainekos, H. Kress-Gazit, and G. Pappas, "Temporal logic motion planning for mobile robots," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation*, 2005, pp. 2020–2025.
- [3] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
- [4] M. Lahijanian, S. B. Andersson, and C. Belta, "Temporal logic motion planning and control with probabilistic satisfaction guarantees," *IEEE Transactions on Robotics*, vol. 28, no. 2, pp. 396–409, 2012.
- [5] V. Raman, A. Donzé, M. Maasoumy, R. M. Murray, A. Sangiovanni-Vincentelli, and S. A. Seshia, "Model predictive control with signal temporal logic specifications," in *53rd IEEE Conference on Decision and Control*, 2014, pp. 81–87.
- [6] C.-I. Vasile, J. Tumova, S. Karaman, C. Belta, and D. Rus, "Minimum-violation sclt motion planning for mobility-on-demand," in *2017 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2017, pp. 1481–1488.
- [7] K. Liang and C.-I. Vasile, "Fair planning for mobility-on-demand with temporal logic requests," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1283–1289.
- [8] G. A. Cardona, D. Saldaña, and C.-I. Vasile, "Planning for modular aerial robotic tools with temporal logic constraints," in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 2878–2883.
- [9] A. Rodionova, L. Lindemann, M. Morari, and G. Pappas, "Temporal robustness of temporal logic specifications: Analysis and control design," *ACM Transactions on Embedded Computing Systems*, vol. 22, no. 1, pp. 1–44, 2022.
- [10] V. Kurtz and H. Lin, "A more scalable mixed-integer encoding for metric temporal logic," *IEEE Control Systems Letters*, vol. 6, pp. 1718–1723, 2021.
- [11] J. Tumova, S. Karaman, C. Belta, and D. Rus, "Least-violating planning in road networks from temporal logic specifications," in *2016 ACM/IEEE 7th International Conference on Cyber-Physical Systems (ICCP)*. IEEE, 2016, pp. 1–9.
- [12] A. Donzé and O. Maler, "Robust satisfaction of temporal logic over real-valued signals," in *Proceedings of the International Conference on Formal Modeling and Analysis of Timed Systems*, 2010.
- [13] M. L. Puterman, *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [14] L. Gurobi Optimization, "Gurobi optimizer reference manual," 2020. [Online]. Available: <http://www.gurobi.com>
- [15] R. Alami, A. Clodic, V. Montreuil, E. A. Sisbot, and R. Chatila, "Toward human-aware robot task planning," in *AAAI spring symposium: to boldly go where no human-robot team has gone before*, 2006, pp. 39–46.
- [16] M. Patel and S. Chernova, "Proactive robot assistance via spatio-temporal object modeling," *arXiv preprint arXiv:2211.15501*, 2022.