

Minimum-Violation Planning for Autonomous Systems: Theoretical and Practical Considerations

Tichakorn Wongpiromsarn¹, Konstantin Slutsky¹, Emilio Frazzoli² and Ufuk Topcu³

Abstract— This paper considers the problem of computing an optimal trajectory for an autonomous system that is subject to a set of potentially conflicting rules. First, we introduce the concept of prioritized safety specifications, where each rule is expressed as a temporal logic formula with its associated weight and priority. The optimality is defined based on the violation of such prioritized safety specifications. We then introduce a class of temporal logic formulas called si-FLTL_{G_X} and develop an efficient, incremental sampling-based approach to solve this minimum-violation planning problem with guarantees on asymptotic optimality. We illustrate the application of the proposed approach in autonomous vehicles, showing that si-FLTL_{G_X} formulas are sufficiently expressive to describe many traffic rules. Finally, we discuss practical considerations and present simulation results for a vehicle overtaking scenario.

I. INTRODUCTION

Autonomous vehicles are subject to several road rules. Often, these rules cannot be simultaneously satisfied. For example, item 221 of Singapore’s *Final Theory of Driving* [1] suggests keeping a safe gap of one meter when passing by a parked vehicle, while item 52 of Singapore’s *Basic Theory of Driving* [2] prohibits crossing a solid double white lane divider. As a result, when encountering a vehicle that is improperly parked in a lane with a solid double white lane divider as shown in Figure 1, an autonomous vehicle may need to violate either of the aforementioned rules unless the lane is wide enough to laterally accommodate two cars with a buffer of one meter.

Previous work shows that linear temporal logic (LTL) [3] is a powerful language for specifying complex properties such as traffic rules [4]–[6]. Existing controller synthesis algorithms for LTL specifications include closed system synthesis [7], reactive synthesis [5], [8], [9], and probabilistic synthesis [10]–[12]. Closed system synthesis considers a deterministic system and generates a controller, if one exists, for the system to satisfy the specification. In contrast, reactive synthesis considers a system operating in a nondeterministic, adversarial environment and ensures that the system satisfies the specification for all possible adversarial actions. Finally, probabilistic synthesis considers a probabilistic system and maximizes the probability that the system satisfies the specification.

*This work was partially supported by AFC Robotics Center of Excellence with award number W911NF1920333.

¹Tichakorn Wongpiromsarn and Konstantin Slutsky are with Iowa State University, IA nok@iastate.edu, k Slutsky@gmail.com

²Emilio Frazzoli is with ETH Zurich, Switzerland efrazzoli@ethz.ch

³Ufuk Topcu is with the University of Texas at Austin, TX utopcu@utexas.edu

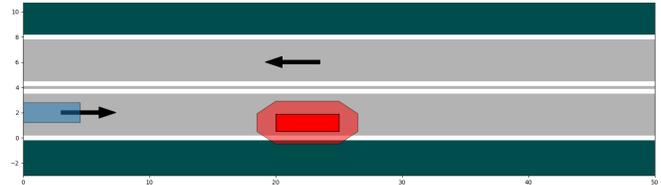


Fig. 1. The autonomous vehicle (blue rectangle) encounters a stationary vehicle (red rectangle) on a two-lane road with a double white lane divider. The red octagon represents the clearance zone around the stationary vehicle.

More recently, minimum-violation planning has been proposed to handle conflicting objectives [13]–[15]. As opposed to reactive and probabilistic synthesis, minimum-violation planning considers a deterministic system and relies on real-time re-planning to respond to quickly changing environments. It is particularly suitable for applications such as autonomous vehicles, where (1) it is hard to obtain an accurate probabilistic model of the environment, (2) accounting for all possible adversarial actions of the environment may render the system too conservative, and (3) the system is subject to multiple rules of different importance and there may be situations where not all the rules can be simultaneously satisfied. Existing work on minimum-violation planning, however, relies on converting an LTL specification to a finite automaton, whose size is exponential in the length of the specification, making real-time re-planning unrealistic.

As shown in [6], most rules of the road can be expressed by a safety formula. In particular, this paper considers the case where autonomous vehicles need to reach the target location but may violate some road rules if needed. We assume that each rule has a certain penalty associated with its violation. The goal of motion planning is to minimize such penalties. The main focus of the paper is on reducing the computational complexity of minimum-violation planning to make real-time re-planning possible.

This paper is closely related to the framework proposed in [16]. The main focus of [16] is the framework for describing rules and their associated violation penalties, taking into various considerations, including law, ethics, local driving culture, etc. Given all the rules and their associated violation penalties, this paper focuses on computing a trajectory of an autonomous vehicle that minimizes the total penalty.

The main contributions of this paper are twofold. First, we introduce a class of linear temporal logic formulas called si-FLTL_{G_X} that is sufficient to precisely describe many traffic rules. Second, we propose an efficient algorithm for computing a path that minimizes the amount of rule violation with the same computational complexity as traditional motion planning algorithms. The proposed algorithm eliminates the

exponential part of the complexity of existing algorithms by avoiding the conversion of specifications into finite automata. The remainder of the paper is organized as follows: Section II introduces the terminology and notations used throughout the paper. Section III formulates the minimum-violation planning problem. Section IV and Section V discuss the solution and practical considerations, respectively. Finally, Section VI presents simulation results.

II. PRELIMINARIES

We consider time-invariant dynamical systems and use finite linear temporal logic (FLTL) [17] to specify their correct behaviors. Given a natural number n , let $\mathbb{N}_{\leq n} = \{0, 1, \dots, n\}$ be the set of natural numbers not greater than n . For any set S , $|S|$ and 2^S denote the cardinality and the powerset of S , respectively.

A. Systems

Let $S \subset \mathbb{R}^d$, where $d \in \mathbb{N}$, be a compact set of states and U be a compact set of control signals. Additionally, let Π denote a finite set of atomic propositions that capture the properties of interest. We define the labeling function $\mathcal{L} : S \rightarrow 2^\Pi$, which maps each state to a set of atomic propositions that are true at that state.

Consider a time-invariant dynamical system

$$\dot{x}(t) = f(x(t), u(t)), \quad (1)$$

with the initial state $x(0) = s_{\text{init}} \in S$. $f : S \times U \rightarrow \mathbb{R}^d$ is assumed to be Lipschitz continuous in both arguments.

Given $T \in \mathbb{R}_{\geq 0}$, $x : [0, T] \rightarrow S$ is a *trajectory* of (1) if there exists $u : [0, T] \rightarrow U$ satisfying (1) for all $t \in [0, T]$. Note that a trajectory does not necessarily start at s_{init} . Let $\mathcal{F}(x) = T$ denote the final time of $x : [0, T] \rightarrow S$.

Consider a trajectory $x : [0, T] \rightarrow S$. A *finite timed word* of x with respect to a finite set $\Gamma = \{t_1, t_2, \dots, t_n\}$ with $0 < t_1 < t_2 < \dots < t_n < T$ is a finite sequence $\omega(x, \Gamma) = l_0^t l_1^t \dots l_n^t$ where $l_i^t = (\lim_{t \rightarrow t_i^+} \mathcal{L}(x(t)), t_{i+1} - t_i)$ for all $i \in \mathbb{N}_{\leq n}$, $t_0 = 0$ and $t_{n+1} = T$. Let $\mathcal{T}(x) = \{t \in (0, T) \mid \lim_{t' \rightarrow t^-} \mathcal{L}(x(t')) \neq \lim_{t' \rightarrow t^+} \mathcal{L}(x(t'))\}$ be the set of discontinuities of $\mathcal{L}(x(\cdot))$. Throughout the paper, we assume that the labeling function \mathcal{L} is such that the limits in the definition of $\mathcal{T}(x)$ exist and $\mathcal{T}(x)$ is finite, for any trajectory x of (1). The *finite timed word* of x is defined as $\omega(x) = \omega(x, \mathcal{T}(x))$.

A dynamical system in (1) can be abstracted into a finite state system represented by a durational Kripke Structure.

Definition 1 (Durational Kripke Structure): A *durational Kripke structure* is a tuple

$$\mathcal{K} = (S_{\mathcal{K}}, s_{\text{init}, \mathcal{K}}, \mathcal{R}_{\mathcal{K}}, \Pi_{\mathcal{K}}, \mathcal{L}_{\mathcal{K}}, \Delta_{\mathcal{K}}), \quad (2)$$

where $S_{\mathcal{K}}$ is a finite set of states, $s_{\text{init}, \mathcal{K}} \in S_{\mathcal{K}}$ is the initial state, $\mathcal{R}_{\mathcal{K}} \subseteq S_{\mathcal{K}} \times S_{\mathcal{K}}$ is a transition relation, $\Pi_{\mathcal{K}}$ is a set of atomic propositions, $\mathcal{L}_{\mathcal{K}} : S_{\mathcal{K}} \rightarrow 2^{\Pi_{\mathcal{K}}}$ is a state labeling function, and $\Delta_{\mathcal{K}} : \mathcal{R}_{\mathcal{K}} \rightarrow \mathbb{R}_{\geq 0}$ is a function assigning a time duration to each transition.

A *finite trace* of \mathcal{K} is a finite sequence of states $\tau = s_0 s_1 \dots s_n$ such that $s_0 = s_{\text{init}, \mathcal{K}}$ and $(s_i, s_{i+1}) \in \mathcal{R}_{\mathcal{K}}$, for all $i \in \mathbb{N}_{\leq n-1}$.

B. Specifications

We consider specifications that can be described by a subset of FLTL. Roughly, an FLTL formula is built up from (a) a set of atomic propositions, (b) the logic connectives: negation (\neg), disjunction (\vee), conjunction (\wedge) and material implication (\implies), and (c) the temporal operators: next (\circ), always (\square), eventually (\diamond) and until (\mathcal{U}). We refer to an FLTL formula that does not include temporal operators as a *propositional logic* formula.

An FLTL formula φ over Π is interpreted over a finite word $w = l_0 l_1 \dots l_n \in (2^\Pi)^{n+1}$ and we write $w \models \varphi$ if w satisfies φ . In particular, consider $p, p' \in \Pi$. $w \models p$ if and only if $p \in l_0$. $w \models \square p$ if and only if $p \in l_i$ for all $i \in \mathbb{N}_{\leq n}$. Consider a more complicated specification that will be used throughout the paper

$$\varphi = \square(p \implies (\circ p \vee \circ p')). \quad (3)$$

In this case, $w \models \varphi$ if and only if for any $i \in \mathbb{N}_{\leq n-1}$ such that $p \in l_i$, we have $p \in l_{i+1}$ or $p' \in l_{i+1}$. The satisfaction of an FLTL formula with respect to a finite timed word can be defined in a natural way: Given an FLTL formula φ over Π and a finite timed word $\omega = (l_0, d_0) \dots (l_n, d_n) \in (2^\Pi \times \mathbb{R}_{\geq 0})^{n+1}$, $\omega \models \varphi$ if and only if $w(\omega) = l_0 \dots l_n \models \varphi$.

Definition 2 (si-FLTL_{G_X}): An si-FLTL_{G_X} formula over a set Π of atomic propositions is an FLTL formula that is stutter-invariant (see below) and is of the form

$$\varphi = \square P_X, \quad (4)$$

where P_X belongs to the smallest set defined inductively by the following rules:

- p is a formula for all $p \in \Pi \cup \{\text{True}, \text{False}\}$,
- $\circ p$ is a formula for all $p \in \Pi \cup \{\text{True}, \text{False}\}$, and
- if P_X^1 and P_X^2 are formulas, then so are $\neg P_X^1$, $P_X^1 \vee P_X^2$, $P_X^1 \wedge P_X^2$ and $P_X^1 \implies P_X^2$.

In other words, P_X is a Boolean combination of propositions from Π and expressions of the form $\circ p$ where $p \in \Pi$.

Roughly, a specification is stutter-invariant if its satisfaction with respect to any word is not affected by operations that duplicate some letters or remove some duplicate letters in that word. For example, consider $w = l_0 l_1 \dots l_n$ and $w' = l_0 l_1 \dots l_{i-1} l_i l_i l_{i+1} \dots l_n$, which is constructed from w by duplicating l_i for some $i \in \mathbb{N}_{\leq n}$. If φ is stutter-invariant, then $w \models \varphi$ if and only if $w' \models \varphi$. We refer the reader to [18] for the definition of stutter-invariant properties. See, e.g., [19], [20] for approaches to check whether a specification is stutter-invariant.

Regardless of its simplicity, si-FLTL_{G_X} turns out to be sufficiently expressive to describe many traffic rules. Reference [21] shows that all the rules enforced in the DARPA Urban Challenge 2007 can be expressed with si-FLTL_{G_X} formulas. All the traffic rules in the examples presented in [15] can also be described using si-FLTL_{G_X} formulas.

Definition 3 (Prioritized Safety Specification): A *prioritized safety specification* is a tuple $\mathcal{P} = (\Pi, \Phi, \Psi, \rho)$ where Π is a set of atomic propositions, Φ is a set of si-FLTL_{G_X} formulas over Π , $\Psi = (\Psi_0, \Psi_1, \dots, \Psi_N)$ organizes the

formulas in Φ into a hierarchy based on their priorities such that $\Psi_i \subseteq \Phi$, for all $i \in \mathbb{N}_{\leq N}$, and $\rho : \Phi \rightarrow \mathbb{N}$ is a function that assigns the weight to each $\varphi \in \Phi$. Throughout the paper, we refer to each $\varphi \in \Phi$ as an *atomic safety rule*.

We use the *level of unsafety* to measure the violation of an si-FLTL $_{G_X}$ formula. Consider an si-FLTL $_{G_X}$ formula $\varphi = \square P_X$ and a finite timed word $\omega = l_0^t l_1^t \dots l_n^t$ where $l_i^t = (l_i, d_i) \in 2^\Pi \times \mathbb{R}_{\geq 0}$ for all $i \in \mathbb{N}_{\leq n}$. We let $l_{n+1}^t = (l_n, 0)$ and define the level of unsafety of ω with respect to φ as

$$\lambda(\omega, \varphi) = \sum_{i \in \mathbb{N}_{\leq n} \mid l_i l_{i+1} \not\models P_X} \tilde{\lambda}(l_i^t, P_X), \quad (5)$$

where $\tilde{\lambda}(l_i^t, P_X) = d_i$ if $l_i l' \not\models P_X$ for all $l' \in 2^\Pi$; otherwise $\tilde{\lambda}(l_i^t, P_X) = 1$. Note that this choice of $\tilde{\lambda}$ differentiates the violation caused by visiting a (unsafe) state with label l_i (the case where $l_i l' \not\models P_X$ for all $l' \in 2^\Pi$) and the violation caused by taking a (unsafe) transition from a state with label l_i to a state with label l_{i+1} (the case where $l_i l_{i+1} \not\models P_X$ but $l_i l' \models P_X$ for some $l' \in 2^\Pi$). In particular, the cost of visiting an unsafe state is the time spent on that state, whereas the cost of taking an unsafe transition is 1. This choice of violation cost is to better accommodate the notion of next as discussed in Remark 1.

Let $\mathcal{P} = (\Pi, \Phi, \Psi, \rho)$ be a prioritized safety specification where $\Psi = (\Psi_0, \Psi_1, \dots, \Psi_N)$. We define the level of unsafety of ω with respect to \mathcal{P} as

$$\lambda_{\mathcal{P}}(\omega) = (\lambda_{\mathcal{P}}(\omega, \Psi_0), \dots, \lambda_{\mathcal{P}}(\omega, \Psi_N)) \in \mathbb{R}^{N+1}, \quad (6)$$

where for each $i \in \mathbb{N}_{\leq N}$,

$$\lambda_{\mathcal{P}}(\omega, \Psi_i) = \sum_{\varphi \in \Psi_i} \rho(\varphi) \lambda(\omega, \varphi). \quad (7)$$

Remark 1: In [15], the level of unsafety of a finite timed word ω with respect to an atomic safety rule φ is defined as

$$\lambda(\omega, \varphi) = \min_{I \subseteq \mathbb{N}_{\leq n} \mid \text{vanish}(\omega, I) \models \varphi} \sum_{i \in I} d_i, \quad (8)$$

where for any given finite sequence $w = l_0 l_1 \dots l_n$ and a set $I \subseteq \mathbb{N}_{\leq n}$, $\text{vanish}(w, I)$ is defined as a subsequence of w obtained by removing all l_i , $i \in I$. This definition is consistent with (5) for the case where φ is an invariant property, i.e., P_X does not include the next operator. However, for the case where P_X includes the next operator, (5) and (8) may yield different results. Our choice of (5) is to better accommodate the notion of next, which is not handled in [15]. For example, consider a finite timed word $\omega = (\{p_0\}, d_0)(\{p_1\}, d_1)$ and an si-FLTL $_{G_X}$ formula $\varphi = \square(p_0 \implies \bigcirc p_0)$ where $\Pi = \{p_0, p_1\}$. In this case, we get $\sum_{i \in \mathbb{N}_{\leq 1} \mid l_i l_{i+1} \not\models P_X} \tilde{\lambda}(l_i, P_X) = 1$ regardless of the value of d_0, d_1 ; thus, the level of unsafety defined in (5) corresponds to the number of unsafe transitions. As a result, this definition allows us to specify an objective such as minimizing the number of lane changes. In contrast, $\min_{I \subseteq \mathbb{N}_{\leq 1} \mid \text{vanish}(\omega, I) \models \varphi} \sum_{i \in I} d_i = \min(d_0, d_1)$, and thus, the level of unsafety defined in (8) corresponds to the total duration either before or after the unsafe transition.

III. PROBLEM FORMULATION

Consider the time-invariant dynamical system (1) with the initial state s_{init} and the set $S_{\text{goal}} \subset S$ of goal states. Let $X = \{x : [0, T] \rightarrow S \mid T \in \mathbb{R}_{\geq 0}, x(0) = s_{\text{init}}, x(T) \in S_{\text{goal}}\}$ be the set of trajectories of (1), starting at s_{init} and ending at a state $s \in S_{\text{goal}}$.

Given a prioritized safety specification $\mathcal{P} = (\Pi, \Phi, \Psi, \rho)$, the minimum-violation planning problem is to compute an optimal trajectory $x^* \in X$ that minimizes the time of reaching a state $s \in S_{\text{goal}}$ among all the trajectories that minimize the level of unsafety with respect to \mathcal{P} . Formally, we define the cost function $J : X \rightarrow \mathbb{R}^{N+2}$ as

$$J(x) = (\lambda_{\mathcal{P}}(\omega(x)), \mathcal{F}(x)). \quad (9)$$

Recall from Section II that $\mathcal{F}(x)$ denotes the final time of trajectory x . As a result, the last coordinate of the cost function J corresponds to the minimum-time objective whereas the first coordinate, $\lambda_{\mathcal{P}}(\omega(x))$, corresponds to the level of unsafety of x with respect to \mathcal{P} .

Using the cost function J , we formally define the minimum-violation planning as follows.

Problem 1 (Minimum-Violation Planning): Based on the standard lexicographical order, compute an optimal trajectory $x^* = \arg \min_{x \in X} J(x)$.

Remark 2: We choose the minimum-time objective as indicated by the maneuver cost function \mathcal{F} for the simplicity of the presentation. Our approach also applies to other maneuver costs, including the control effort, with some minor modifications.

IV. SOLUTION

Reference [15] solves Problem 1 by constructing a weighted finite automaton \mathcal{A} that is the product of weighted finite automata, each corresponding to an atomic safety rule $\varphi \in \Phi$. The weights on the transitions of \mathcal{A} are defined such that the weight of the shortest accepting run over any word ω is the level of unsafety of ω . The product $\mathcal{K} \otimes \mathcal{A}$ of the Kripke structure \mathcal{K} and \mathcal{A} is incrementally constructed. It can be shown that Problem 1 is equivalent to finding a shortest path in $\mathcal{K} \otimes \mathcal{A}$.

As the size of \mathcal{A} is exponential in the length of φ [3], our approach avoids constructing the product $\mathcal{K} \otimes \mathcal{A}$ to reduce computational complexity. Instead, we translate an si-FLTL $_{G_X}$ formula over Π into an si-FLTL $_G$ formula over $\Pi \times \Pi$. As will be discussed later, this translation allows us to incrementally construct and maintain only the Kripke structure \mathcal{K} (as opposed to $\mathcal{K} \otimes \mathcal{A}$ as in the aforementioned work), and compute the weights of its transitions based on the satisfaction of propositional formulas of the consecutive states and the time duration of the transitions. As a result, it allows temporal logic specifications to be handled with the same computational complexity as traditional motion planning algorithms such as RRT* and RRG.

Definition 4 (si-FLTL $_G$): An si-FLTL $_G$ formula over $\Pi \times \Pi$ is an si-FLTL $_{G_X}$ formula $\varphi = \square P$ where P is a propositional logic formula over $\Pi \times \Pi$.

A propositional logic formula P over $\Pi \times \Pi$ is interpreted over a pair $(l, l') \in 2^\Pi \times 2^\Pi$ with the satisfaction relation

\models defined as follows. For $p, p' \in \Pi \cup \{\text{True}, \text{False}\}$ and $(l, l') \in 2^\Pi \times 2^\Pi$, $(l, l') \models (p, p')$ if and only if $l \models p$ and $l' \models p'$. Here, for any $l \in 2^\Pi$, we have $l \models \text{True}$, $l \not\models \text{False}$, and for any $p \in \Pi$, $l \models p$ if and only if $p \in l$. The logic connectives are defined as in the standard propositional logic.

Based on the semantics of FLTL, given a finite word $w = l_0 l_1 \dots l_n \in (2^\Pi)^{n+1}$ and an si-FLTL_G formula $\varphi = \square P$ over $\Pi \times \Pi$, we say that w satisfies φ , written $w \models_{\Pi \times \Pi} \varphi$ if and only if $(l_i, l_{i+1}) \models P$ for all $i \in \mathbb{N}_{\leq n-1}$ and $(l_n, l_n) \models P$. Note that the terminal condition $(l_n, l_n) \models P$ results from the assumption that φ is stutter-invariant, which ensures that $w \models \varphi$ if and only if $w' = l_0 l_1 \dots l_n l_n \models \varphi$.

The level of unsafety of a finite timed word $\omega = l_0^t l_1^t \dots l_n^t$ with respect to an si-FLTL_G formula φ over $\Pi \times \Pi$ is defined by

$$\lambda(\omega, \varphi) = \sum_{i \in \mathbb{N}_{\leq n} \mid (l_i, l_{i+1}) \not\models P} \tilde{\lambda}(l_i^t, P), \quad (10)$$

where $l_i^t = (l_i, d_i) \in 2^\Pi \times \mathbb{R}_{\geq 0}$ for all $i \in \mathbb{N}_{\leq n}$, $l_{n+1}^t = (l_n, 0)$, $\tilde{\lambda}(l_i^t, P) = d_i$ if $(l_i, l') \not\models P$ for all $l' \in 2^\Pi$, otherwise $\tilde{\lambda}(l_i^t, P) = 1$.

A. Conversion of si-FLTL_{G_X} to si-FLTL_G

Given an si-FLTL_{G_X} formula φ over Π , we define an operation denext that constructs an si-FLTL_G formula over $\Pi \times \Pi$ from φ by replacing each instance of p in φ with (p, True) and replacing each instance of $\circ p$ in φ with (True, p) for all $p \in \Pi$. For example, consider an si-FLTL_{G_X} formula φ defined in (3). The corresponding si-FLTL_G formula over $\Pi \times \Pi$ is given by

$$\text{denext}(\varphi) = \square \left((p, \text{True}) \implies ((\text{True}, p) \vee (\text{True}, p')) \right). \quad (11)$$

Lemma 1: Let $\Pi = \{p_0, p_1, \dots, p_n\}$ and $\Pi' = \{q_0, q_1, \dots, q_{n'}\}$ be sets of propositions. Let P be a propositional formula over $\Pi \cup \Pi'$ and P' be a propositional formula over $\Pi \times \Pi'$ that is constructed from P by replacing each instance of p_i by (p_i, True) and replacing each instance of q_j by (True, q_j) for all $i \in \mathbb{N}_{\leq n}$, $j \in \mathbb{N}_{\leq n'}$. Then, for any $l \subseteq \Pi$ and $l' \subseteq \Pi'$, $l \cup l' \models P$ if and only if $(l, l') \models P'$.

Proof: Consider arbitrary $l \subseteq \Pi$ and $l' \subseteq \Pi'$, $i \in \mathbb{N}_{\leq n}$ and $j \in \mathbb{N}_{\leq n'}$. It follows directly from the definition of True that (a) $l \models p_i$ if and only if $(l, l') \models (p_i, \text{True})$, and (b) $l' \models q_j$ if and only if $(l, l') \models (\text{True}, q_j)$. As a result, we can conclude from the construction of P' and the semantics of propositional logic that $l \cup l' \models P$ if and only if $(l, l') \models P'$. ■

Lemma 2: Let P_X be defined as in Definition 2. Consider a propositional logic formula P' over $\Pi \times \Pi$ that is constructed from P_X by replacing each instance of p in P_X with (p, True) and replacing each instance of $\circ p$ in P_X with (True, p) for all $p \in \Pi \cup \{\text{True}, \text{False}\}$. For any arbitrary $l, l' \in 2^\Pi$, we have $l \cup l' \models P_X$ if and only if $(l, l') \models P'$.

Proof: Let $\Pi = \{p_0, p_1, \dots, p_n\}$. Define $\Pi' = \{q_0, q_1, \dots, q_n\}$. Let P be a propositional formula over $\Pi \cup \Pi'$ that is constructed from P_X by replacing each instance of $\circ p_i$ by q_i for all $i \in \mathbb{N}_{\leq n}$.

Consider arbitrary $l, l' \in 2^\Pi$. Let $\Pi_{l'} = \{q_i \mid p_i \in l'\} \subseteq \Pi'$. Based on the definition of the \circ operator, we can conclude that $l \cup l' \models P_X$ if and only if $l \cup \Pi_{l'} \models P$.

By construction, P' is obtained from P by replacing each instance of p_i in P by (p_i, True) and replacing each instance of q_i in P by (True, p_i) for all $i \in \mathbb{N}_{\leq n}$. As a result, we can conclude using Lemma 1 and the definition of $\Pi_{l'}$ that $l \cup \Pi_{l'} \models P$ if and only if $(l, l') \models P'$. Combining this with the result from the previous paragraph, we obtain $l \cup l' \models P_X$ if and only if $(l, l') \models P'$. ■

We now establish the equivalence of the level of unsafety with respect to an si-FLTL_{G_X} formula over Π and the level of unsafety with respect to the corresponding si-FLTL_G formula over $\Pi \times \Pi$.

Lemma 3: For any finite timed word ω and any si-FLTL_{G_X} formula φ over Π ,

$$\lambda(\omega, \varphi) = \lambda(\omega, \text{denext}(\varphi)). \quad (12)$$

Proof: This result can be trivially derived from Lemma 2 and the definitions of the level of unsafety (5) and (10). ■

Finally, we construct the prioritized safety specification $\hat{\mathcal{P}} = (\Pi \times \Pi, \hat{\Phi}, \hat{\Psi}, \hat{\rho})$ with each atomic safety rule obtained from that of \mathcal{P} by applying denext operation. Formally, $\hat{\Phi} = \{\text{denext}(\varphi) \mid \varphi \in \Phi\}$, $\hat{\Psi} = (\hat{\Psi}_0, \hat{\Psi}_1, \dots, \hat{\Psi}_N)$, $\hat{\Psi}_i = \{\text{denext}(\varphi) \mid \varphi \in \Psi_i\}$ for all $i \in \mathbb{N}_{\leq N}$, and $\hat{\rho}(\varphi) = \hat{\rho}(\text{denext}(\varphi))$ for all $\varphi \in \Phi$. The level of unsafety of a finite timed word ω with respect to $\hat{\mathcal{P}}$ is defined following (6), (7) as $\lambda_{\hat{\mathcal{P}}}(\omega) = (\lambda_{\hat{\mathcal{P}}}(\omega, \hat{\Psi}_0), \dots, \lambda_{\hat{\mathcal{P}}}(\omega, \hat{\Psi}_N)) \in \mathbb{R}^{N+1}$, where for each $i \in \mathbb{N}_{\leq N}$, $\lambda_{\hat{\mathcal{P}}}(\omega, \hat{\Psi}_i) = \sum_{\varphi \in \hat{\Psi}_i} \hat{\rho}(\varphi) \lambda(\omega, \varphi)$. Based on the construction of $\hat{\mathcal{P}}$ and Lemma 3, we obtain the following result, which allows us to replace \mathcal{P} with $\hat{\mathcal{P}}$.

Proposition 1: For any finite timed word ω , $\lambda_{\mathcal{P}}(\omega) = \lambda_{\hat{\mathcal{P}}}(\omega)$.

B. Incremental Construction of Weighted Kripke Structure

We follow a sampling-based procedure described in [22] to incrementally construct a Kripke structure \mathcal{K} as a finite state representation of the dynamical system (1). The main difference is that we augment \mathcal{K} with weights on its transitions. The weights are picked such that the sum of the weights on any finite trace τ of \mathcal{K} is the level of unsafety of the finite timed word generated by the trajectory of (1) corresponding to τ . As opposed to [15], we do not construct the weighted product automaton $\mathcal{K} \otimes \mathcal{A}$ where \mathcal{A} is created by combining all the automata, each corresponding to each $\varphi \in \Phi$.

Definition 5 (Weighted Kripke Structure): A *weighted Kripke structure* is a tuple

$$\bar{\mathcal{K}} = (S_{\mathcal{K}}, s_{\text{init}, \mathcal{K}}, \mathcal{R}_{\mathcal{K}}, \Pi_{\mathcal{K}}, \mathcal{L}_{\mathcal{K}}, \mathcal{W}_{\mathcal{K}}), \quad (13)$$

where $S_{\mathcal{K}}$, $s_{\text{init}, \mathcal{K}}$, $\mathcal{R}_{\mathcal{K}}$, $\Pi_{\mathcal{K}}$ and $\mathcal{L}_{\mathcal{K}}$ are defined as in Definition 1 and $\mathcal{W}_{\mathcal{K}} : \mathcal{R}_{\mathcal{K}} \rightarrow \mathbb{R}_{\geq 0}^n$ for some $n \in \mathbb{N}$ is a function assigning a transition cost to each transition in $\mathcal{R}_{\mathcal{K}}$.

A *finite trace* of $\bar{\mathcal{K}}$ is defined as that of \mathcal{K} . Given a finite trace $\tau = s_0 s_1 \dots s_n$, we define the weight of τ as $\mathcal{W}(\tau) = \sum_{i \in \mathbb{N}_{\leq n-1}} \mathcal{W}_{\mathcal{K}}(s_i, s_{i+1})$. For any $s \in S_{\mathcal{K}}$,

define $\text{Traces}(\bar{\mathcal{K}}, s) = \{\tau = s_0 s_1 \dots s_n \mid n \in \mathbb{N}, s_n = s, \tau \text{ is a finite trace of } \bar{\mathcal{K}}\}$ to be the set of all the finite traces of $\bar{\mathcal{K}}$ that end at s .

Various sampling-based algorithms such as RRT*, RRG, and their k -nearest variants can be employed to incrementally construct $\bar{\mathcal{K}}$ [22]. The key difference between these algorithms lies in the connections of states. In particular, the RRT* algorithm maintains a tree structure rather than a graph as in RRG, ensuring that each state only has at most one parent. It maintains an upper bound on the cost $J_{\mathcal{K}}(s)$ of the unique path from the initial state to each state $s \in S_{\mathcal{K}}$. Algorithm 1 provides a common template for incrementally constructing $\bar{\mathcal{K}}$, based on the following primitive procedures.

- a) *Sample*: $\text{sample} : \mathbb{N} \rightarrow S$ is a function that generates independent, identically distributed samples from a uniform distribution supported over S .
- b) *Add*: Given a state s , $\text{add}(s)$ adds s to $S_{\mathcal{K}}$, i.e., it executes $S_{\mathcal{K}} \leftarrow S_{\mathcal{K}} \cup \{s\}$. For RRT*, it also sets $J_{\mathcal{K}}(s) = 0$ if $s = s_{\text{init}}$; otherwise $J_{\mathcal{K}}(s) = \infty$.
- c) *Steer*: Given states $s, s' \in S$, $\text{steer}(s, s')$ returns the set of trajectories $x : [0, T] \rightarrow S$ of (1) such that $x(0) = s$, $x(T) = s'$, and $\mathcal{T}(x)$ exists and is finite.
- d) *Nearest neighbors*: $\text{near} : S \rightarrow 2^{S_{\mathcal{K}}}$ computes the set of nearest neighbors. When applying the RRT* or RRG algorithm, we let

$$\text{near}(s) = \left\{ s' \in S_{\mathcal{K}} \mid \|s' - s\|_2 \leq \left(\gamma \frac{\log m}{m} \right)^{1/D} \right\},$$

where m is the cardinality of $S_{\mathcal{K}}$, and γ and D are constants that depend on the dimension d of the state space and the Lebesgue measure of S . For the case of k -nearest RRG or k -nearest RRT*, $\text{near}(s)$ returns k nearest neighbors of $s \in S$ where $k > \gamma' \log(n)$ for some constant γ' . We refer the reader to [22], [23] for the definitions of γ , γ' , and D .

- e) *Transition cost*: For any trajectory $x : [0, T] \rightarrow S$ of (1), we define the cost of x as $\mathcal{C}(x) = (\lambda_{\hat{p}}(\omega(x)), T) \in \mathbb{R}^{N+2}$. Note that according to (9) and Proposition 1, $\mathcal{C}(\cdot)$ corresponds to the cost function defined in Problem 1. Let $\omega(x) = (l_0, d_0)(l_1, d_1) \dots (l_n, d_n)$. Thanks to (10), $\lambda_{\hat{p}}(\omega(x))$ can be computed by simply evaluating whether (l_i, l_{i+1}) satisfies the propositional formula corresponding to each atomic safety rule for each $i \in \mathbb{N}_{\leq n}$, with $l_{n+1} = l_n$. For any states $s, s' \in S$, the transition cost from s to s' is then defined for the case where $\text{steer}(s, s') \neq \emptyset$ as $\text{cost}(s, s') = \min_{x \in \text{steer}(s, s')} \mathcal{C}(x)$.
- f) *Connect*: Given states $s, s' \subseteq S_{\mathcal{K}}$, $\text{connect}(s, s')$ updates the relevant elements based on a transition from s to s' . For RRG, $\text{connect}(s, s')$ simply adds the transition (s, s') to $\mathcal{R}_{\mathcal{K}}$ and set $\mathcal{W}_{\mathcal{K}}(s, s') = \text{cost}(s, s')$ as shown in Algorithm 2. In contrast, as shown in Algorithm 3, the transition (s, s') is added for RRT* only if it improves the cost to reach s' from s_{init} (Line 1). If so, existing transitions to s' are removed (Line 2) and $J_{\mathcal{K}}(s')$ is updated (Line 4). By propagating the change in $J_{\mathcal{K}}(s')$ down the tree structure, we obtain the RRT# algorithm, which ensures that the *promising* vertices (i.e., those that

have the potential to be part of the optimal solution) are consistent, i.e., $J_{\mathcal{K}}(\tilde{s})$ is the cost of the unique path from the initial state to a promising vertex $\tilde{s} \in S_{\mathcal{K}}$. This approach has been shown to improve the convergence rate of RRT*. See [24] for more details on RRT#.

Algorithm 1 returns the weighted Kripke structure $\bar{\mathcal{K}}_n$ after n iterations as well as the set $S_{\text{goal}, \mathcal{K}} = S_{\text{goal}} \cap S_{\mathcal{K}}$ of the sampled goal states. For each state $s \in S_{\mathcal{K}}$, let $\tau(\bar{\mathcal{K}}_n, s) \in \text{Traces}(\bar{\mathcal{K}}_n, s)$ denote an optimal trace of $\bar{\mathcal{K}}_n$ that ends at s , i.e., $\mathcal{W}(\tau(\bar{\mathcal{K}}_n, s)) \leq \mathcal{W}(\tilde{\tau})$ for all $\tilde{\tau} \in \text{Traces}(\bar{\mathcal{K}}_n, s)$. We define $s_{\text{goal}, \mathcal{K}} = \arg \min_{s \in S_{\text{goal}, \mathcal{K}}} \mathcal{W}(\tau(\bar{\mathcal{K}}_n, s))$. Note that for the case of RRT#, $s_{\text{goal}, \mathcal{K}} = \arg \min_{s \in S_{\text{goal}, \mathcal{K}}} J_{\mathcal{K}}(s)$ and for any promising vertex $s \in S_{\mathcal{K}}$, $\mathcal{W}(\tau(\bar{\mathcal{K}}_n, s)) = J_{\mathcal{K}}(s)$ and $\tau(\bar{\mathcal{K}}_n, s)$ can be obtained by following the unique parent of each state backward, starting from s to the initial state, i.e., $\tau(\bar{\mathcal{K}}_n, s) = s_0 s_1 \dots s_m$ for some $m \in \mathbb{N}$ such that $s_0 = s_{\text{init}}$, $s_m = s$, and $s_i = \text{parent}(s_{i+1})$ for all $i \in \mathbb{N}_{\leq m-1}$. Here, $\text{parent}(\tilde{s})$ is a unique state with $(\text{parent}(\tilde{s}), \tilde{s}) \in \mathcal{R}_{\mathcal{K}}$. When the connections are constructed based on the RRT* or RRG algorithm, $s_{\text{goal}, \mathcal{K}}$ and $\tau(\bar{\mathcal{K}}_n, s)$ can be obtained using, e.g., the Dijkstra's shortest path algorithm.

Algorithm 1: Minimum-violation planning.

```

1  $S_{\mathcal{K}} \leftarrow \emptyset; \mathcal{R}_{\mathcal{K}} \leftarrow \emptyset; S_{\text{goal}, \mathcal{K}} \leftarrow \emptyset;$ 
2  $\text{add}(s_{\text{init}});$ 
3 foreach  $i \in \mathbb{N}_{\leq n}$  do
4    $s_{\text{new}} \leftarrow \text{sample}(i);$ 
5    $S_{\text{near}} = \text{near}(s_{\text{new}});$ 
6    $\text{add}(s_{\text{new}});$ 
7   foreach  $s \in S_{\text{near}}$  do
8     if  $\text{steer}(s, s_{\text{new}}) \neq \emptyset$  then
9        $\text{connect}(s, s_{\text{new}});$ 
10  foreach  $s \in S_{\text{near}}$  do
11    if  $\text{steer}(s_{\text{new}}, s) \neq \emptyset$  then
12       $\text{connect}(s_{\text{new}}, s);$ 
13  if  $s_{\text{new}} \in S_{\text{goal}}$  then
14     $S_{\text{goal}, \mathcal{K}} = S_{\text{goal}, \mathcal{K}} \cup \{s_{\text{new}}\};$ 
15 return  $\bar{\mathcal{K}}_n = (S_{\mathcal{K}}, s_{\text{init}}, \mathcal{R}_{\mathcal{K}}, \Pi, \mathcal{L}, \mathcal{W}_{\mathcal{K}}), S_{\text{goal}, \mathcal{K}}$ 

```

Algorithm 2: $\text{connect}(s, s')$ for RRG.

```

1  $\mathcal{R}_{\mathcal{K}} \leftarrow \mathcal{R}_{\mathcal{K}} \cup \{(s, s')\};$ 
2  $\mathcal{W}_{\mathcal{K}}(s, s') = \text{cost}(s, s')$ 

```

Algorithm 3: $\text{connect}(s, s')$ for RRT*.

```

1 if  $J_{\mathcal{K}}(s) + \text{cost}(s, s') < J_{\mathcal{K}}(s')$  then
2    $\mathcal{R}_{\mathcal{K}} \leftarrow (\mathcal{R}_{\mathcal{K}} \setminus \{(s_1, s_2) \in \mathcal{R}_{\mathcal{K}} \mid s_2 = s'\}) \cup \{(s, s')\};$ 
3    $\mathcal{W}_{\mathcal{K}}(s, s') = \text{cost}(s, s');$ 
4    $J_{\mathcal{K}}(s') = J_{\mathcal{K}}(s) + \text{cost}(s, s');$ 

```

Let $c^* = \min_{s \in S_{\text{goal}}} \text{cost}(s_{\text{init}}, s)$ be the cost of an optimal trajectory of (1) from s_{init} to S_{goal} based on the cost function

C. The following result can be directly derived from the asymptotic optimality of the RRG and RRT* algorithms [22] and the boundedness of cost for the case where P_X in (4) is a propositional formula for all $\varphi \in \Phi$.

Lemma 4: If P_X in (4) is a propositional formula for all $\varphi \in \Phi$, then $\mathcal{W}(\tau(\bar{\mathcal{K}}_n, s_{\text{goal}, \mathcal{K}}))$ converges to c^* almost surely, i.e.,

$$\mathbb{P}\left(\left\{\lim_{n \rightarrow \infty} \mathcal{W}(\tau(\bar{\mathcal{K}}_n, s_{\text{goal}, \mathcal{K}})) = c^*\right\}\right) = 1.$$

We now establish the equivalence of the transition cost cost and the original cost function J of Problem (1).

Lemma 5: $c^* = \min_{x \in X} J(x)$.

Proof: This result follows directly from Proposition 1 and the definitions of J and cost. ■

Consider an arbitrary $s \in S_{\mathcal{K}}$. Let $\tau(\bar{\mathcal{K}}_n, s) = s_0 s_1 \dots s_m$ for some $m \in \mathbb{N}$. For each $i \in \mathbb{N}_{\leq m-1}$, let $x_{s,i} = \arg \min_{x \in \text{steer}(s_i, s_{i+1})} \mathcal{C}(x)$, i.e., $x_{s,i}$ is a trajectory of (1) from s_i to s_{i+1} with $\mathcal{C}(x_{s,i}) = \text{cost}(s_i, s_{i+1})$. A trajectory x_s of (1) from s_{init} to s can be constructed from $\tau(\bar{\mathcal{K}}_n, s)$ by concatenating $x_{s,i}$. Formally, $x_s : [0, \sum_{i \in \mathbb{N}_{\leq m-1}} T_i] \rightarrow S$ such that $x_s(t + \sum_{k=0}^{i-1} T_k) = x_{s,i}(t)$ for all $t \in [0, T_i]$ and $i \in \mathbb{N}_{\leq m-1}$ where $T_i = \mathcal{F}(x_{s,i})$.

Lemma 6: For any $s \in S_{\mathcal{K}}$, $\mathcal{W}(\tau(\bar{\mathcal{K}}_n, s)) = J(x_s)$.

Proof: Let $\tau(\bar{\mathcal{K}}_n, s) = s_0 s_1 \dots s_m$. By construction, $\mathcal{W}(\tau(\bar{\mathcal{K}}_n, s)) = \sum_{i \in \mathbb{N}_{\leq m-1}} \text{cost}(s_i, s_{i+1})$. Additionally, from the definition of \mathcal{C} , J and Proposition 1, $J(x_s) = \sum_{i \in \mathbb{N}_{\leq m-1}} \mathcal{C}(x_{s,i})$. We can then conclude from the definition of $x_{s,i}$ that $J(x_s) = \sum_{i \in \mathbb{N}_{\leq m-1}} \text{cost}(s_i, s_{i+1})$. ■

Combining Lemma 4–6, we obtain the asymptotic optimality of Algorithm 1.

Proposition 2: Let x^* be a solution of Problem 1 and $J^* = J(x^*)$. If P_X in (4) is a propositional formula for all $\varphi \in \Phi$, then $J(x_{s_{\text{goal}, \mathcal{K}}})$ converges to J^* almost surely.

If there exists $\varphi \in \Phi$ such that P_X is not a propositional formula, then the asymptotic optimality of Algorithm 1 cannot be guaranteed as cost is not necessarily bounded. However, Lemma 6 ensures that $\tau(\bar{\mathcal{K}}_n, s_{\text{goal}, \mathcal{K}})$ is an optimal trajectory among those in $\bar{\mathcal{K}}_n$.

The analysis in [15] shows that the computational complexity of the original Minimum-Violation RRT* algorithm is $\mathcal{O}(K^2 n \log n)$ where K is the number of states in the weighted finite automaton \mathcal{A} , which is exponential in the length of the specification. As Algorithm 1 is a special case of the RRT* and RRG algorithms where the transition cost corresponds to the level of unsafety of the transition, it shares the same computational complexity of $\mathcal{O}(n \log n)$ as that of the RRT* and RRG algorithms, which is the same as that of the original Minimum-Violation RRT* algorithm with $K = 1$. As in [15], this analysis relies on the assumption that the complexity of $\text{connect}(s, s')$ is $\mathcal{O}(1)$ for all $s, s' \in S_{\mathcal{K}}$.

V. PRACTICAL CONSIDERATIONS

The asymptotic optimality of Algorithm 1 is essential in many safety-critical applications as it ensures that a sufficiently safe trajectory will be found, if one exists,

given sufficient computation time. For autonomous vehicles, however, the available computation time is often limited due to the dynamic nature of the environments in which they operate. Such environments include not only relatively static features such as road markings, constructions, weather conditions, etc., but also dynamic features arising from other agents (vehicles, pedestrians, animals, etc.) sharing the road. An optimal trajectory with respect to the environment at time $t \in \mathbb{R}_{\geq 0}$ may become the least safe option with respect to the environment at time $t + \epsilon$, even for small $\epsilon \in \mathbb{R}_{\geq 0}$, especially when there is a drastic change in the environment. Examples of such situations include (a) a newly detected object shows up, (b) another agent violates the right of way, and (c) a vehicle that is initially parked starts to move off while the autonomous vehicle is overtaking it. In fact, a similar situation to the latter led to an accident between the Cornell and the MIT autonomous vehicles during the 2007 DARPA Urban Challenge [25].

Reactive synthesis [5], [8], [9] and probabilistic synthesis [10]–[12] have been applied to handle dynamic environments. Roughly, in reactive synthesis, a control policy is constructed to ensure that the system satisfies its specification for all valid environment behaviors. In contrast, probabilistic synthesis considers a probabilistic model of the environment and constructs a control policy that maximizes the probability that the system satisfies its specification. Both approaches assume a good understanding of the environment: Reactive synthesis requires the knowledge of all possible behaviors of the environment, whereas probabilistic synthesis relies on an accurate probabilistic model of the environment. The control protocol synthesized by these approaches could be invalid in that the system could be left with no valid trajectory if such assumptions do not hold.

A key advantage of minimum-violation planning compared to these approaches is in handling unexpected or unmodeled environment behaviors: As long as there exist trajectories from s_{init} to a state in S_{goal} in $\bar{\mathcal{K}}_n$, the algorithm always returns the safest one. Hence, the richness of $\bar{\mathcal{K}}_n$ is a crucial factor for successful applications of this approach.

As autonomous vehicles are required to respond quickly to changes in the environment, $\bar{\mathcal{K}}_n$ needs to be updated in each planning iteration. To ensure the richness of $\bar{\mathcal{K}}_n$, we maintain $\bar{\mathcal{K}}_n$ from the previous planning iteration instead of rebuilding it from scratch. Then, each planning iteration updates s_{init} and $\mathcal{W}_{\mathcal{K}}$ based on the current state of the vehicle and the most recently observed environment. Additionally, due to changes in $\mathcal{W}_{\mathcal{K}}$, RRG-based algorithms are potentially more preferable than RRT*-based algorithms.

In summary, we initialize $\bar{\mathcal{K}}_n$ with $S_{\mathcal{K}} = \{s_{\text{init}}\}$ and $\mathcal{R}_{\mathcal{K}} = \emptyset$. Then, each planning iteration performs the following procedures.

- (i) Update s_{init} and $\mathcal{W}_{\mathcal{K}}$ based on the current state of the system.
- (ii) Augment $\bar{\mathcal{K}}_n$ based on lines 3–14 of Algorithm 1.
- (iii) Extract an optimal trajectory in $\bar{\mathcal{K}}_n$ using a graph search algorithm (e.g., Dijkstra or A*).

Finally, Algorithm 1 can be adapted to other definitions

of the level of unsafety by modifying the transition cost function \mathcal{C} such that $J(x_s) = \sum_{i \in \mathbb{N}_{\leq m-1}} \mathcal{C}(x_{s,i})$ for all $s \in S_{\mathcal{K}}$ where $s_0 s_1 \dots s_m = \tau(\bar{\mathcal{K}}_n, s)$. In this case, it is easy to show that Lemma 5 and Lemma 6 still hold; thus, an optimal trace of $\bar{\mathcal{K}}_n$ to a goal state as extracted by a graph search algorithm still corresponds to a safest trajectory among all the options in $\bar{\mathcal{K}}_n$. In fact, we have experimented with other definitions of the level of unsafety as suggested by the rulebooks framework [16], including non-additive costs [26]. Certain definitions, however, invalidate the assumption of Lemma 4 and as a result, lead to the loss of asymptotic optimality guarantee similar to the case where P_X in (4) is not a propositional formula. Refining the definition of the level of unsafety as well as ensuring the richness of $\bar{\mathcal{K}}_n$ are subject to the current research.

VI. EXPERIMENTAL RESULTS

We consider an autonomous vehicle modeled by a Dubins car [27]: $\dot{x} = \cos(\theta)$, $\dot{y} = \sin(\theta)$ and $\dot{\theta} = u$ where $u \in [-1, 1]$, (x, y) is the position of the center of the rear axle and θ is the heading of the vehicle. The autonomous vehicle encounters a stationary vehicle while it is navigating a two-lane road with a solid yellow center line as show in Figure 1. The set of goal states is given by $S_{\text{goal}} = \{(x, y, \theta) \mid x \geq 37\}$.

A. Prioritized Safety Specification

For any $x, y, \theta \in \mathbb{R}$, let $\text{FP}(x, y, \theta) \subset \mathbb{R}^2$ be the footprint of the autonomous vehicle when the center of its rear axle is at (x, y) and its heading is θ . Consider atomic propositions collision, close, road and lane, representing the autonomous vehicle colliding with the stationary vehicle, overlapping with the clearance zone, being fully on the road and being fully within a correct lane, respectively. Formally, the labeling function \mathcal{L} is defined such that for any $x, y, \theta \in \mathbb{R}$, (a) $\text{collision} \in \mathcal{L}(x, y, \theta)$ iff $\text{FP}(x, y, \theta) \cap \text{SV} \neq \emptyset$ where $\text{SV} \subset \mathbb{R}^2$ is the footprint of the stationary vehicle, (b) $\text{close} \in \mathcal{L}(x, y, \theta)$ iff $\text{FP}(x, y, \theta) \cap \text{CZ} \neq \emptyset$ where $\text{CZ} \subset \mathbb{R}^2$ is the clearance zone around the stationary vehicle, constructed from SV and the required lateral and longitudinal clearance (see Figure 1), (c) $\text{road} \in \mathcal{L}(x, y, \theta)$ iff $\text{FP}(x, y, \theta) \subseteq \text{RD}$ where RD is the road, i.e., the area where a vehicle is allowed to drive, and (d) $\text{lane} \in \mathcal{L}(x, y, \theta)$ iff $\text{FP}(x, y, \theta) \subseteq \text{LN}$ where LN is the right lane, i.e., the lane with the correct travel direction for the autonomous vehicle.

We consider the following atomic safety rules, each of which can be expressed by an si-FLTL $_{G_X}$ formula.

- (i) **No collision:** $\varphi_1 = \square \neg \text{collision}$.
- (ii) **Staying on road:** $\varphi_2 = \square \text{road}$.
- (iii) **Obstacle clearance:** $\varphi_3 = \square \neg \text{close}$.
- (iv) **Lane keeping:** $\varphi_4 = \square \text{lane}$.

The prioritized safety specification $\mathcal{P} = (\Pi, \Phi, \Psi, \rho)$ is defined as $\Pi = \{\text{collision}, \text{close}, \text{road}, \text{lane}\}$, $\Phi = \{\varphi_1, \dots, \varphi_4\}$, $\Psi = \{\{\varphi_1\}, \{\varphi_2\}, \{\varphi_3, \varphi_4\}\}$, and $\rho(\varphi_i) = 1$, for all i .

B. Simulation Results

Algorithm 1 was implemented in TuLiP, a Python-based software toolbox [28] and run on a laptop with Intel Core i7-10710U processor. 40 iterations of Line 3–14 of Algorithm 1 were run with $n = 20$, i.e., 20 states were added in each iteration. Figure 2 shows the optimal traces $\tau(\bar{\mathcal{K}}_n, s_{\text{goal}, \mathcal{K}})$ when the connection is based on the RRT* algorithm. The optimal cost $\mathcal{W}(\tau(\bar{\mathcal{K}}_n, s_{\text{goal}, \mathcal{K}}))$ at the end of the 40th iteration is $(0, 0, 12.7, 37.0)$. The level of unsafety and computation time of each iteration is shown in Figure 3.

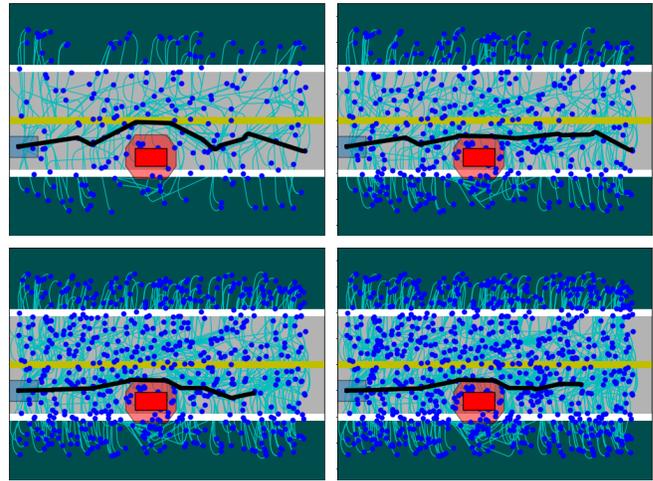


Fig. 2. The states (dark blue dots) and their connections (light blue curves) in $\bar{\mathcal{K}}_n$ and the optimal path (black curve) extracted at the end of the 10th (top left), 20th (top right), 30th (bottom left), and 40th (bottom right) iterations when Algorithm 1 is applied with RRT* connections.

The results when applying RRG connections are shown in Figure 4 and Figure 5. The optimal cost $\mathcal{W}(\tau(\bar{\mathcal{K}}_n, s_{\text{goal}, \mathcal{K}}))$ at the end of the 40th iteration is $(0, 0, 11.9, 35.5)$. The Kripke structure constructed based on the RRG algorithm includes significantly more connections than that constructed based on the RRT* algorithm. For both cases, the level of unsafety with respect to Ψ_1 and Ψ_2 quickly converges to 0.

Note that the majority of the computation time is spent on computing labels along a trajectory. This is similar to the case of traditional motion planning, where collision checking is typically the main bottleneck [27]. In fact, computing the violation of φ_1 is exactly the collision checking problem. The specific implementation in this example performs expensive polygon operations to compute the labels. The computation time can be significantly reduced by employing more efficient polygon operations and parallel computation.

VII. CONCLUSIONS

This paper introduced a class of LTL formulas that are sufficiently expressive to describe traffic rules such as lane-keeping, obstacle avoidance, etc. Given traffic rules specified by these formulas and their relative importance, we proposed an incremental algorithm to compute a trajectory for an autonomous vehicle to reach a given goal while minimizing the level of unsafety with respect to the given rules. Both the theoretical guarantees and practical considerations

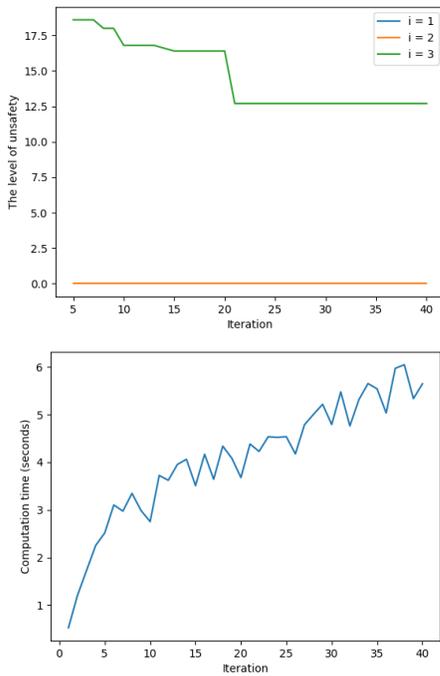


Fig. 3. The level of unsafety of the optimal path with respect to $\Psi_i, i \in \{1, 2, 3\}$ and the computation time (seconds) of each iteration when Algorithm 1 is applied with RRT* connections.

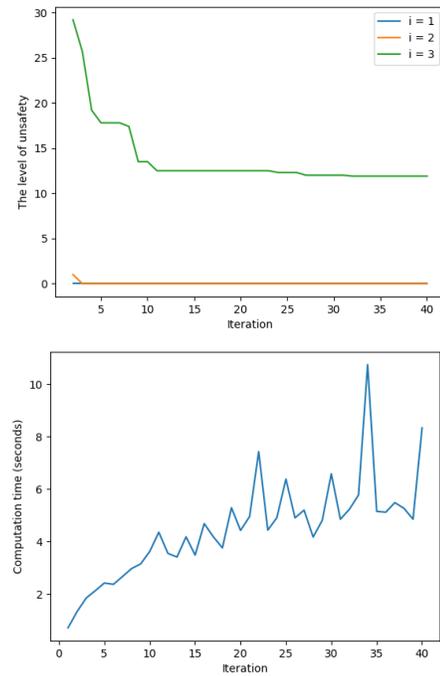


Fig. 5. The level of unsafety of the optimal path with respect to $\Psi_i, i \in \{1, 2, 3\}$ and computation time (seconds) of each iteration when Algorithm 1 is applied with RRG connections.

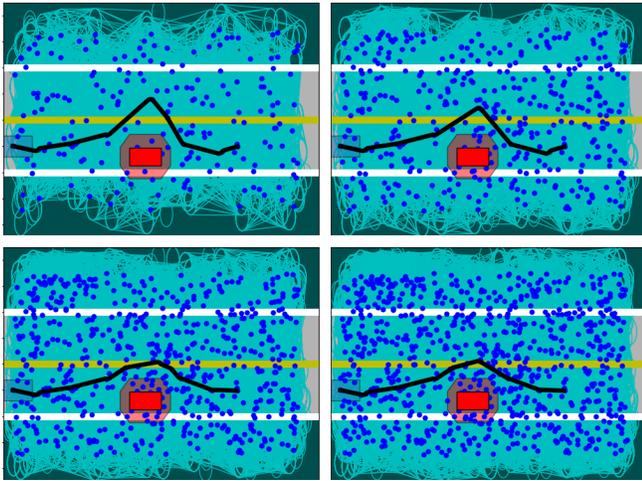


Fig. 4. The states (dark blue dots) and their connections (light blue curves) in $\bar{\mathcal{K}}_n$ and the optimal path (black curve) extracted at the end of the 10th (top left), 20th (top right), 30th (bottom left), and 40th (bottom right) iterations when Algorithm 1 is applied with RRG connections.

were discussed. Simulation results for the vehicle overtaking scenario were provided.

ACKNOWLEDGMENT

The authors gratefully acknowledge Dmitry Yershov for insightful discussions.

REFERENCES

- [1] *Final Theory of Driving: The Official Handbook*, 9th ed. Singapore Traffic Police, July 2017.
- [2] *Basic Theory of Driving: The Official Handbook*, 10th ed. Singapore Traffic Police, February 2018.
- [3] C. Baier and J.-P. Katoen, *Principles of Model Checking (Representation and Mind Series)*. The MIT Press, 2008.

- [4] H. Kress-Gazit, D. Conner, H. Choset, A. Rizzi, and G. Pappas, "Courteous cars: Decentralized multiagent trafficcoordination," *IEEE Robotics Automation Magazine*, vol. 15, no. 1, pp. 30–38, 2008.
- [5] H. Kress-Gazit and G. J. Pappas, "Automatically synthesizing a planning and control subsystem for the DARPA Urban Challenge," in *2008 IEEE International Conference on Automation Science and Engineering*, Aug 2008, pp. 766–771.
- [6] T. Wongpiromsarn, S. Karaman, and E. Frazzoli, "Synthesis of provably correct controllers for autonomous vehicles in urban environments," in *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, Oct 2011, pp. 1168–1173.
- [7] M. Kloetzer and C. Belta, "A fully automated framework for control of linear systems from temporal logic specifications," *IEEE Transactions on Automatic Control*, vol. 53, no. 1, pp. 287–297, 2008.
- [8] T. Wongpiromsarn, U. Topcu, and R. M. Murray, "Receding horizon temporal logic planning," *IEEE Transactions on Automatic Control*, vol. 57, no. 11, pp. 2817–2830, Nov 2012.
- [9] H. Kress-Gazit, T. Wongpiromsarn, and U. Topcu, "Correct, reactive, high-level robot control," *IEEE Robotics Automation Magazine*, vol. 18, no. 3, pp. 65–74, Sep. 2011.
- [10] T. Wongpiromsarn and E. Frazzoli, "Control of probabilistic systems under dynamic, partially known environments with temporal logic specifications," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, Dec 2012, pp. 7644–7651.
- [11] A. Ulusoy, T. Wongpiromsarn, and C. Belta, "Incremental controller synthesis in probabilistic environments with temporal logic constraints," *International Journal of Robotics Research*, vol. 33, no. 8, pp. 1130–1144, Jul 2014.
- [12] 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, April 2012.
- [13] J. Tumova, L. I. R. Castro, S. Karaman, E. Frazzoli, and D. Rus, "Minimum-violation LTL planning with conflicting specifications," in *2013 American Control Conference*, June 2013, pp. 200–205.
- [14] J. Tumova, G. C. Hall, S. Karaman, E. Frazzoli, and D. Rus, "Least-violating control strategy synthesis with safety rules," in *Proceedings of the 16th International Conference on Hybrid Systems: Computation and Control*, 2013, pp. 1–10.
- [15] L. I. R. Castro, P. Chaudhari, J. Tumova, S. Karaman, E. Frazzoli, and D. Rus, "Incremental sampling-based algorithm for minimum-violation motion planning," in *52nd IEEE Conference on Decision and Control*, Dec 2013, pp. 3217–3224.

- [16] A. Censi, K. Slutsky, T. Wongpiromsarn, D. Yershov, S. Pendleton, J. Fu, and E. Frazzoli, "Liability, ethics, and culture-aware behavior specification using rulebooks," in *2019 International Conference on Robotics and Automation (ICRA)*, May 2019, pp. 8536–8542.
- [17] Z. Manna and A. Pnueli, *Temporal Verification of Reactive Systems: Safety*. Berlin, Heidelberg: Springer-Verlag, 1995.
- [18] D. Peled and T. Wilke, "Stutter-invariant temporal properties are expressible without the next-time operator," *Information Processing Letters*, vol. 63, no. 5, pp. 243–246, 1997.
- [19] J. Klein and C. Baier, "On-the-fly stuttering in the construction of deterministic ω -automata," in *Implementation and Application of Automata*, J. Holub and J. Zdárek, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 51–61.
- [20] T. Michaud and A. Duret-Lutz, "Practical stutter-invariance checks for ω -regular languages," in *Proceedings of the 22Nd International Symposium on Model Checking Software - Volume 9232*, ser. SPIN 2015. Berlin, Heidelberg: Springer-Verlag, 2015, pp. 84–101.
- [21] T. Wongpiromsarn, S. Karaman, and E. Frazzoli, "Synthesis of provably correct controllers for autonomous vehicles in urban environments," in *2011 14th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, 2011, pp. 1168–1173.
- [22] S. Karaman and E. Frazzoli, "Sampling-based algorithms for optimal motion planning," *International Journal of Robotics Research*, vol. 30, no. 7, pp. 846–894, June 2011.
- [23] K. Solovey, L. Janson, E. Schmerling, E. Frazzoli, and M. Pavone, "Revisiting the asymptotic optimality of RRT*," in *2020 International Conference on Robotics and Automation (ICRA)*, 2020.
- [24] O. Arslan and P. Tsiotras, "Use of relaxation methods in sampling-based algorithms for optimal motion planning," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 2421–2428.
- [25] L. Fletcher, S. Teller, E. Olson, D. Moore, Y. Kuwata, J. How, J. Leonard, I. Miller, M. Campbell, D. Huttenlocher, A. Nathan, and F.-R. Kline, "The mit–cornell collision and why it happened," *Journal of Field Robotics - Special Issue on the 2007 DARPA Urban Challenge, Part III*, vol. 25, no. 10, pp. 775–807, Oct. 2008.
- [26] K. Slutsky, D. Yershov, T. Wongpiromsarn, and E. Frazzoli, "Optimal path problem and iterated A* algorithm," in *The 14th International Workshop on the Algorithmic Foundations of Robotics (WAFR)*, 2020.
- [27] S. M. LaValle, *Planning Algorithms*. Cambridge University Press, 2006.
- [28] T. Wongpiromsarn, U. Topcu, N. Ozay, H. Xu, and R. M. Murray, "Tulip: A software toolbox for receding horizon temporal logic planning," in *Proceedings of the 14th International Conference on Hybrid Systems: Computation and Control*, 2011, pp. 313–314. [Online]. Available: <https://github.com/tulip-control>