





STLts-Div: Diversified Trace Synthesis from STL Specifications Using MILP

Martin Jouve-Genty^{1,2} , Han Su² , Sota Sato^{2,5} , Jie An³ ,
Zhenya Zhang^{2,4} , and Ichiro Hasuo^{2,5,6}  

¹ Département Informatique, ENS de Lyon, Lyon, France
suhan@nii.ac.jp

² National Institute of Informatics, Tokyo, Japan
sotasato@nii.ac.jp, anjie@iscas.ac.cn

³ Institute of Software, Chinese Academy of Sciences, Beijing, China
zhang@ait.kyushu-u.ac.jp

⁴ Kyushu University, Fukuoka, Japan
hasuo@nii.ac.jp

⁵ SOKENDAI (The Graduate University for Advanced Studies), Tokyo, Japan

⁶ Imiron Co., Ltd, Tokyo, Japan

Abstract. Modern cyber-physical systems are complex, and requirements are often written in Signal Temporal Logic (STL). Writing the right STL is difficult in practice; engineers benefit from concrete executions that illustrate what a specification actually admits. Trace synthesis addresses this need, but a single witness rarely suffices to understand intent or explore edge cases—diverse satisfying behaviors are far more informative. We introduce diversified trace synthesis: the automatic generation of sets of behaviorally diverse traces that satisfy a given STL formula. Building on a MILP encoding of STL and system model, we formalize three complementary diversification objectives—Boolean distance, random Boolean distance, and value distance—all captured by an objective function and solved iteratively. We implement these ideas in STLts-Div, a lightweight Python tool that integrates with Gurobi.

Keywords: Cyber-Physical Systems · Signal Temporal Logic · Trace Synthesis

1 Introduction

Cyber-physical systems (CPS) underpin safety-critical applications such as aircraft collision-avoidance protocols [23], aerospace [2], and pacemakers [24].

M. Jouve-Genty, H. Su: Equal contribution.

The authors are supported by JST ASPIRE Grant No. JPMJAP2301. Jie An is supported by the CAS Project for Young Scientists in Basic Research Grant No. YSBR-123. Zhenya Zhang is supported by JST BOOST Grant No. JPMJBY24D7 and JSPS Grant No. JP25K21179. The work was done during M.J.’s internship at National Inst. of Inform., Japan.

© The Author(s) 2026

A. Sampaio and M. Stoelinga (Eds.): FM 2026, LNCS 16556, pp. 3–20, 2026.

https://doi.org/10.1007/978-3-032-26204-2_1

Ensuring their reliability is essential, yet formal verification is challenging due to hybrid/nonlinear dynamics, continuous state spaces, and environmental uncertainty. As a result, full verification is often infeasible in practice.

Trace synthesis offers a pragmatic alternative: instead of proving properties for all behaviors, it constructs concrete executions that satisfy (or violate) a given temporal specification, providing actionable artifacts for testing and analysis. As a dual to model checking [20], trace synthesis has been realized via SMT encodings [3], optimization-based falsification [8], and—more recently—mixed-integer linear programming (MILP) encodings that capture both *Signal Temporal Logic (STL)* constraints and plant dynamics [20]. These approaches generate informative scenarios and counterexamples while interfacing naturally with simulators and models.

However, a single satisfying trace is often insufficient—for example, when validating whether an STL requirement truly captures the intent of the engineer. A *set* of behaviorally diverse traces that illuminate different ways the specification can be met is required. In this work we study *diversified trace synthesis* from STL, aiming to generate traces that satisfy the specification while exhibiting sufficient diversity.

We address this problem with STLts-Div, a tool that synthesizes sets of diverse traces for a given system model \mathcal{M} and STL specification φ . STLts-Div builds on the MILP-based synthesis framework of [20]. We introduce three complementary diversity objectives—Boolean distance (BD), randomized Boolean distance (RBD), and value distance (VD)—and encode them as MILP objectives. We adopt the MILP setting because its numeric decision variables allow these objectives to be integrated directly, whereas propositional SMT encodings do not readily support summing/counting of truth assignments (see Remark 2).

Contributions and Organization. We summarize our contributions as below.

- We formulate the problem of diversified trace synthesis for STL.
- We propose three diversity objectives (BD, RBD, VD) and provide MILP objective function encodings compatible with [20].
- We implement these ideas in STLts-Div, a Python package, and evaluate them on benchmark suites, demonstrating substantially higher diversity than a solution-pool baseline.

Section 2 reviews STL preliminaries and states the problem. Section 3 recaps the MILP encoding of [20], which underpins our approach. Section 4 formalizes three diversity objectives and presents our MILP-based synthesis procedures. Section 5 demonstrates the STLts-Div tool. Section 6 reports experiments evaluating effectiveness and efficiency. We conclude in Sect. 7.

Related Work. Signal Temporal Logic is widely used for specifying CPS behaviors. Prior work spans controller synthesis [6, 17, 18, 21], runtime/online monitoring [22, 25], and trace synthesis/falsification [3, 20], among others. The choice of technique depends on the task: for continuous-time control, Lyapunov/barrier-based methods synthesize controllers that enforce STL constraints [13]; for bounded model checking, SMT encodings search for a time partition witnessing

satisfaction [3]; and recent MILP encodings support trace synthesis by jointly searching over signal values and a δ -stable partition (see Sect. 3) [20]. Existing approaches, however, typically return one satisfying execution. In contrast, we target the multi-trace setting and synthesize sets of behaviorally diverse traces for a fixed model and STL formula.

2 Preliminaries

Let \mathbb{N} , \mathbb{R} , and $\mathbb{R}_{\geq 0}$ denote the sets of natural numbers, real numbers, and non-negative real numbers, respectively. Let $\mathbb{B} = \{\top, \perp\}$ denote the Boolean domain. For a set X , we write $\wp(X)$ for its power set and $|X|$ for its cardinality. An *interval* is any subset of $\mathbb{R}_{\geq 0}$ of the form (a, b) , $(a, b]$, $[a, b)$, or $[a, c]$, where $a < b$ and $a \leq c$.

Signal. Let $T > 0$ and V be a finite set of variables. A *signal* over V with time horizon T is a function $\sigma: [0, T] \rightarrow \mathbb{R}^{|V|}$. We let \mathbf{Signal}_V^T denote the set of all such signals; when T is clear from the context we write \mathbf{Signal}_V . For $v \in V$ we write $\sigma(t)(v)$ for the value of the variable v in σ at time t . We extend any $\sigma \in \mathbf{Signal}_V^T$ to $t > T$ by setting $\sigma(t) := \sigma(T)$ for all $t > T$.

For a signal σ and $t \geq 0$, the *t-postfix* is the signal σ^t defined by $\sigma^t(t') := \sigma(t + t')$. The *t-postfix* is the basis for defining the STL semantics (see [11]).

System model. Let V' and V be the input and output variable sets. A *system model* \mathcal{M} from V' to V with a time horizon T is a (possibly nondeterministic) map $\mathcal{M}: \mathbf{Signal}_{V'}^T \rightarrow \wp(\mathbf{Signal}_V^T)$. The trace set of \mathcal{M} is $\mathcal{L}(\mathcal{M}) := \bigcup_{\tau \in \mathbf{Signal}_{V'}^T} \mathcal{M}(\tau)$, i.e., the set of all output signals under any input τ .

Example 1 (Single-lane car). Consider a car moving in a single lane whose dynamics are given by the ordinary differential equation (ODE)

$$\dot{x} = v, \quad \dot{v} = a, \tag{1}$$

where x , v and a denote position, velocity and acceleration, respectively.

We formalize this system as \mathcal{M}_{car} with input variables $V' = \{a, v^{\text{init}}, x^{\text{init}}\}$ and output variables $V = \{a, v, x\}$. Here, the acceleration a changes over time; the initial velocity and position $v^{\text{init}}, x^{\text{init}}$ change as well but only their initial values $\tau(0)(v^{\text{init}}), \tau(0)(x^{\text{init}})$ matter. Given an input signal τ , the model yields $\mathcal{M}_{\text{car}}(\tau) = \{\sigma\}$, where σ is determined by Eq. (1): $\sigma(t)(a) = \tau(t)(a)$, $\sigma(t)(v) = \tau(0)(v^{\text{init}}) + \int_0^t \tau(\theta)(a) d\theta$, $\sigma(t)(x) = \tau(0)(x^{\text{init}}) + \int_0^t \sigma(\theta)(v) d\theta$.

2.1 Signal Temporal Logic (STL)

We use the standard syntax and semantics of STL [14, 20]. We present some basic definitions for the record.

In STL, an *atomic proposition* over a variable set V is represented as $p := (\pi_p(\vec{w}) \geq 0)$, where $\pi_p: \mathbb{R}^{|V|} \rightarrow \mathbb{R}$ maps a valuation $\vec{w} \in \mathbb{R}^{|V|}$ of variables in

V to a real value. In this work, we restrict attention to *linear* (affine) atomic propositions, i.e., $\pi_p(\vec{w}) = a^\top \vec{w} + c$ with $a \in \mathbb{R}^{|V|}$ and $c \in \mathbb{R}$. This syntactic restriction enables an MILP encoding of STL formulas.

Any STL formula has an equivalent *negation normal form* (NNF) [9], in which negations appear only adjacent to atomic propositions. We assume all STL formulas are in NNF in this work. The syntax of STL is as follows:

$$\varphi ::= \top \mid p \mid \neg p \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \varphi_1 \mathcal{U}_I \varphi_2 \mid \varphi_1 \mathcal{R}_I \varphi_2,$$

where I is a nonsingular closed time interval, and \mathcal{U}_I and \mathcal{R}_I are temporal operators *until* and *release*, respectively. The *eventually* and *always* operators are syntactic sugars defined as $\diamond_I \varphi = \top \mathcal{U}_I \varphi$, $\square_I \varphi = \perp \mathcal{R}_I \varphi$. The set of all subformulas of φ is denoted by $\text{Sub}(\varphi)$.

The Boolean semantics $\sigma \models \varphi$ and robust semantics $\llbracket \sigma, \varphi \rrbracket \in \mathbb{R} \cup \{+\infty, -\infty\}$ of STL are standard. Details can be found in the Appendix A of [11].

2.2 Finite Variability

Finite variability is a standard assumption used to make STL model checking and synthesis tractable in practice (see, e.g., [12, 15]). In particular, without such a restriction the full (unbounded) STL model-checking problem is EXPSPACE-complete [1]; adopting finite variability permits efficient bounded analyses, which we also assume in this work.

Definition 1 (finite variability [16]). *A Boolean signal $q : \mathbb{R}_{\geq 0} \rightarrow \mathbb{B}$ has N -bounded variability if there exists a collection $\mathcal{P} = \{J_i\}_{i=1}^N$ of nonempty, pairwise disjoint intervals with $\bigcup_{i=1}^N J_i = \mathbb{R}_{\geq 0}$ such that q is constant on every $J_i \in \mathcal{P}$.*

Intuitively, N -bounded variability means that the Boolean signal can change its value at most $N - 1$ times over $\mathbb{R}_{\geq 0}$. We say that a signal σ has N -bounded variability *with respect to* an STL formula φ if the Boolean signal $t \mapsto (\sigma^t \models \varphi)$ has the N -bounded variability. Moreover, σ has *hereditary N -bounded variability* (abbreviated N -HBV) with respect to φ if σ has N -bounded variability with respect to every subformula $\psi \in \text{Sub}(\varphi)$.

We recall the bounded STL trace synthesis problem from [20]:

Problem 1. Given an STL formula φ (over V), a system model \mathcal{M} (from V' to V) with time horizon T , and a variability bound $N \in \mathbb{N}$, find a trace $\sigma \in \mathcal{L}(\mathcal{M})$ such that i) σ has the N -HBV with respect to φ and ii) $\sigma \models \varphi$ holds, or prove such σ does not exist.

In this paper we study a *diversity-aware* variant of the above problem: instead of returning an arbitrary satisfying trace, we generate a set of traces that are mutually diverse, i.e., sufficiently distinct under a suitable notion of difference.

Formalizing diversity is nontrivial: it can be quantified in multiple, fundamentally different ways. In Sect. 4 we discuss this issue and present three distinct methods for measuring and enforcing diversity in MILP-based synthesis.

3 MILP-Based Trace Synthesis

We recall the MILP-based trace synthesis workflow in [20]. The key idea, originally from [3], is to *partition* the time horizon into time intervals, with respect to the truth values of relevant STL formulas. The partition should be such that

- it ensures *interval-wise constancy*: the truth value of each subformula $\psi \in \text{Sub}(\varphi)$ remains constant on each interval;
- it is *MILP-encodable*: it avoids strict inequalities ($<$), since MILP solvers cannot distinguish $<$ from \leq and only support non-strict ones (\leq) [10].

The original SMT-based approach [3] partitions the horizon into alternating singleton and open intervals, i.e., $\dots, (\gamma_{i-1}, \gamma_i), \{\gamma_i\}, (\gamma_i, \gamma_{i+1}), \dots$. However, this scheme requires $<$ to distinguish open and closed intervals and is therefore not MILP-encodable. To address this, we use closed intervals $\dots, [\gamma_{i-1}, \gamma_i], [\gamma_i, \gamma_{i+1}], \dots$ to partition the horizon.

Closed intervals can be expressed only using \leq ; but they have overlaps (at γ_i). This overlap creates ambiguity for the previous notion of interval-wise constancy, which requires that a subformula is either “constantly true” or “never true” on an interval. The technique of *δ -stable partition* [20] avoids this ambiguity via introducing the *δ -tightening* of a STL formula with some $\delta > 0$ (Def. 4).

Definition 2 (time sequence, time state sequence [20]). A time sequence over $[0, T]$ is a sequence $\Gamma = (0 = \gamma_0 < \dots < \gamma_N = T)$. The same construct is often denoted by $\Gamma = ([\gamma_{i-1}, \gamma_i])_{i=1}^N$, using the corresponding partition of $[0, T]$ into closed intervals with singular overlaps. We write Γ_i for $[\gamma_{i-1}, \gamma_i]$.

Given a time sequence, a time state sequence over V is a sequence $\varsigma = ((x_0, \gamma_0), \dots, (x_N, \gamma_N))$, where x_0, \dots, x_N in $\mathbb{R}^{|V|}$.

Definition 3 (piecewise-linear signal [20]). Given a timed state sequence $\varsigma = ((x_0, \gamma_0), \dots, (x_N, \gamma_N))$, the piecewise-linear signal $\varsigma^{\text{pwl}}: [0, \gamma_N] \rightarrow \mathbb{R}^{|V|}$ is defined by the following linear interpolation: $\varsigma^{\text{pwl}} := (1 - \lambda)x_{i-1} + \lambda x_i$ if $\gamma_{i-1} \leq t \leq \gamma_i$ (where $\lambda = \frac{1}{\gamma_i - \gamma_{i-1}}(t - \gamma_{i-1})$).

Definition 4 (δ -tightening of STL formulas in NNF [20]). Let φ be an STL formula in NNF whose atomic propositions are linear predicates of the form $(c^\top x + b \geq 0)$. Given $\delta > 0$, the δ -tightening φ^δ is obtained from φ by i) replacing every occurrence of $(c^\top x + b \geq 0)$ with $(c^\top x + b \geq \delta)$, and ii) replacing every occurrence of $\neg(c^\top x + b \geq 0)$ with $(-c^\top x - b \geq \delta)$. All Boolean connectives and temporal connectives (with the time intervals therein) are left unchanged.

Note that φ^δ is stronger than φ , i.e., $(\sigma \models \varphi^\delta) \Rightarrow (\sigma \models \varphi)$ for all signals σ . Using this notion, a δ -stable partition defined below classifies each subformula on each closed interval as either i) always true, or ii) never *robustly* true.

Definition 5 (δ -stability [20]). Let φ be an STL formula over V in NNF, $\sigma \in \text{Signal}_V^T$ be a signal, and $\Gamma = (\gamma_0, \dots, \gamma_N)$ be a time sequence (Def. 2) with $\gamma_N = T$. We say Γ is δ -stable for σ and φ if, for each $i = 1, \dots, N$ and each subformula $\psi \in \text{Sub}(\varphi)$, either of the following holds: i) $\sigma^t \models \psi$ for each $t \in \Gamma_i$, or ii) $\sigma^t \not\models \psi^\delta$ for each $t \in \Gamma_i$.

Finding a δ -stable partition Γ for a signal σ and formula φ , so that $\sigma^t \models \varphi$ holds for $t \in \Gamma_1$, is a sufficient method to prove $\sigma \models \varphi$. Recall that $\Gamma_1 = [\gamma_0, \gamma_1]$ with $\gamma_0 = 0$ (see Def. 2). As shown in [20], this search can be encoded as an MILP when σ is piecewise-linear (Def. 3).

Definition 6 (conservative valuation [20]). *Let φ be an STL formula in NNF, and $\Gamma = (\gamma_0, \dots, \gamma_N)$ be a time sequence. A valuation of φ in Γ is a function $\Theta: \text{Sub}(\varphi) \times \{1, \dots, N\} \rightarrow \mathbb{B}$ that assigns, to each subformula and index of the intervals of Γ , a Boolean truth value. Let σ be a signal with a time horizon $T = \gamma_N$. We say that Θ is a conservative valuation of φ in Γ on σ (up to δ) if i) $\Theta(\psi, i) = \top$ implies that, for each $t \in \Gamma_i$, $\sigma^t \models \psi$ holds; and ii) $\Theta(\psi, i) = \perp$ implies, for each $t \in \Gamma_i$, $\sigma^t \not\models \psi^\delta$.*

Lemma 1 ([20]). *Suppose there is a conservative valuation Θ of an STL formula φ in a time sequence Γ on σ up to δ . Then Γ is δ -stable for σ and φ .*

We simply write $\langle \psi \rangle_i$ for $\Theta(\psi, i)$ when Θ is clear from context. The MILP problem therefore searches for i) a partition $\Gamma = ([\gamma_{i-1}, \gamma_i])_{i=1}^N$ (see Def. 2), and ii) a conservative valuation Θ (see Def. 6), such that:

- the truth values assigned to subformulas comply with the STL semantics (Sect. 2.1), e.g., for a subformula $(\psi_1 \wedge \psi_2) \in \text{Sub}(\varphi)$ and $i \in \{1, \dots, N\}$, $\langle \psi_1 \wedge \psi_2 \rangle_i = \langle \psi_1 \rangle_i \wedge \langle \psi_2 \rangle_i$;
- Θ assigns \top to the top-level formula φ in Γ_1 ; and
- there is a piecewise-linear trace $\sigma \in \mathcal{L}(\mathcal{M})$ of \mathcal{M} that yields Θ .

The entities Γ, Θ we search for are expressed as MILP variables, and the above three conditions are expressed as MILP constraints. The variables can be classified into four types, where $N \in \mathbb{N}$ is a constant for variability bound (Prob. 1).

- Real-valued variables $\gamma_0, \dots, \gamma_N$ specifying the time sequence Γ .
- Boolean variables $\{\langle \psi \rangle_i \mid 1 \leq i \leq N, \psi \in \text{Sub}(\varphi)\}$ for the value $\Theta(\psi, i)$ of a valuation Θ that we search for.
- Real-valued variables $\{x_{i,v} \mid 0 \leq i \leq N, v \in V\}$ for the values of a piecewise-linear trace $\sigma \in \mathcal{L}(\mathcal{M})$.
- Auxiliary variables (real or binary) used to encode big- M constraints—a common technique in LP—and the temporal/satisfaction encodings.

We omit the MILP details here; see [20] for the full encoding.

4 Diversity-Aware Synthesis

This section presents methods to synthesize mutually diverse traces. Although modern MILP solvers such as Gurobi [10] can return multiple solutions from the solution pool of a single instance, those solutions are typically very similar (see Sect. 6 for an empirical demonstration).

To overcome this limitation, we propose three diversity measures that quantify how “different” one trace is from another. Each measure is encoded as an MILP objective to promote diverse solutions. We do not claim that the three methods presented here are the only possible approaches. However, given the practical need to generate diverse example traces, it is important to provide methods that work.

4.1 Boolean Distance Method (BD)

As flipping the truth value of a subformula may induce substantial changes in the underlying signal variables, an initial attempt to diversify traces is to bias the truth valuations of subformulas in each trace. We quantify the resulting diversity using a Boolean distance between traces.

Definition 7 (Boolean distance). *Let σ_1, σ_2 be two traces found by an MILP solver that i) have N -HBV (see Prob. 1) and ii) satisfy the same STL formula φ .*

Let Γ^{σ_1} (resp. Γ^{σ_2}) be the δ -stable partition for σ_1 (resp. σ_2) and φ . Let $\langle \psi \rangle_i^{\sigma_j}$ denote the Boolean valuation of $\psi \in \text{Sub}(\varphi)$ over the interval $\Gamma_i^{\sigma_j}$ of the partition Γ^{σ_j} , where $i \in \{1, \dots, N\}$ and $j \in \{1, 2\}$. The Boolean distance between σ_1 and σ_2 is

$$d_B(\sigma_1, \sigma_2) := \sum_{i=1}^N \sum_{\psi \in \text{Sub}(\varphi)} |\langle \psi \rangle_i^{\sigma_1} - \langle \psi \rangle_i^{\sigma_2}|. \quad (2)$$

By construction, $d_B(\sigma_1, \sigma_2)$ is equal to the number of pairs (i, ψ) on which the two traces assign different truth values. Hence, a larger Boolean distance indicates greater diversity, in the sense of satisfaction of subformulas.

Remark 1. In Eq. (2), if the Boolean value of $\langle \psi \rangle_i^{\sigma_1}$ is fixed for $i \in \{1, \dots, N\}$ (which is the case in our usage), the absolute value can be rewritten as

$$|\langle \psi \rangle_i^{\sigma_1} - \langle \psi \rangle_i^{\sigma_2}| = \begin{cases} \langle \psi \rangle_i^{\sigma_2} & \text{if } \langle \psi \rangle_i^{\sigma_1} = 0, \\ 1 - \langle \psi \rangle_i^{\sigma_2} & \text{if } \langle \psi \rangle_i^{\sigma_1} = 1, \end{cases}$$

or equivalently, $|\langle \psi \rangle_i^{\sigma_1} - \langle \psi \rangle_i^{\sigma_2}| = (1 - 2\langle \psi \rangle_i^{\sigma_1}) \langle \psi \rangle_i^{\sigma_2} + \langle \psi \rangle_i^{\sigma_1}$.

Note that the last expression is affine in $\langle \psi \rangle_i^{\sigma_2}$ ($\langle \psi \rangle_i^{\sigma_1}$ is a constant). Applying this rule to Eq. (2), we obtain an affine expression over variables $\langle \psi \rangle_i^{\sigma_2}$.

Given a fixed variability bound $N \in \mathbb{N}$, we incorporate the Boolean distance of traces into the MILP objective while keeping the original constraints from Sect. 3 unchanged. This MILP problem is then solved iteratively, for $m = 1, 2, \dots$, to synthesize a trace σ_m that is maximally distinct (in terms of Boolean distance) from the previously synthesized traces $\sigma_1, \dots, \sigma_{m-1}$. The objective in each iteration is defined as follows.

- **Initialization** ($m = 1$). Maximize $f_1 := 0$.
- **Iteration** ($m > 1$). Let σ_m denote the trace to be synthesized at the current iteration (this gives decision variables). Given the previously obtained traces $\{\sigma_i\}_{i=1}^{m-1}$ (which are constants in MILP), maximize

$$f_m := \sum_{i=1}^{m-1} d_B(\sigma_m, \sigma_i). \quad (3)$$

By Remark 1, f_m is an affine function of the Boolean decision variables of the trace σ_m , with coefficients determined by the constants for $\{\sigma_i\}_{i=1}^{m-1}$.

Remark 2. The above diversification method is possible due to our use of MILP. A similar method in SMT is nontrivial. The key difference is that, in MILP, truth values are encoded as integer variables and they can be summed up (as in Eq. (3)), whereas in SMT they are propositional variables and cannot be summed. Counting the number of subformulas with different truth values across models would require techniques such as MaxSMT or model counting, which we leave for future work.

4.2 Randomized Boolean Distance Method (RBD)

This subsection introduces a randomized variant of the Boolean distance. It promotes diversity by evaluating each candidate trace against an independently sampled reference, thereby decoupling the current solution from all previous ones.

Specifically, for a fixed variability bound N , the valuation of all $\psi \in \text{Sub}(\varphi)$ across the intervals $(\Gamma_i)_{i=1}^N$ form an $N|\text{Sub}(\varphi)|$ -dimensional Boolean vector. We sample a *reference valuation* σ_{ref} from the uniform distribution in $\mathbb{B}^{N|\text{Sub}(\varphi)|}$ and solve an MILP problem that maximizes the Boolean distance between σ_{ref} and the candidate trace. Let σ denote the candidate trace, the objective is

$$f := d_B(\sigma, \sigma_{\text{ref}}),$$

which yields a trace that is as far as possible, in Boolean distance, from the reference valuation. By Remark 1, f is affine in the decision variables of trace σ .

Remark 3. The choice between maximization and minimization is immaterial: for any reference valuation $\sigma_{\text{ref}} \in \mathbb{B}^{N|\text{Sub}(\varphi)|}$,

$$d_B(\sigma, \sigma_{\text{ref}}) = N|\text{Sub}(\varphi)| - d_B(\sigma, \bar{\sigma}_{\text{ref}}),$$

where $\bar{\sigma}_{\text{ref}}$ denotes the bitwise complement of σ_{ref} . Hence, maximizing distance from σ_{ref} is equivalent to minimizing distance to $\bar{\sigma}_{\text{ref}}$.

The following proposition provides a quantitative guarantee that the probability of generating identical reference traces is small, which supports the diversity of traces. The proof can be found in Appendix B of [11].

Proposition 1. For a δ -stable partition $\Gamma = (\gamma_0, \dots, \gamma_N)$ of the time horizon for a formula φ , let $d := N |\text{Sub}(\varphi)|$. Suppose $m \geq 2$ reference valuations $\sigma_{\text{ref}}^1, \dots, \sigma_{\text{ref}}^m$ are drawn independently and uniformly from \mathbb{B}^d . Let p denote the probability that at least two references coincide. Then

$$p = 1 - \frac{2^d!}{(2^d - m)! 2^{md}},$$

and the following upper bound holds:

$$p \leq \frac{m(m-1)}{2^{d+1}}.$$

4.3 Value Distance Method (VD)

Although the Boolean distance can serve as a metric for diversification, we also consider a more direct metric: the distance between the trace values. For two piecewise-linear traces that share the same time sequence, their *value distance* is defined as follows.

Definition 8. Let ς_1, ς_2 be two time state sequences (Def. 2) with the same time sequence $\Gamma = (\gamma_0, \dots, \gamma_N)$. Let V be a set of signal variables. For two piecewise-linear traces $\sigma_1, \sigma_2 : [\gamma_0, \gamma_N] \rightarrow \mathbb{R}^{|V|}$ generated from ς_1, ς_2 , respectively, their value distance is defined by

$$d_V(\sigma_1, \sigma_2) := \sum_{i=0}^N \|\sigma_1(\gamma_i) - \sigma_2(\gamma_i)\|_1 = \sum_{i=0}^N \sum_{v \in V} |\sigma_1(\gamma_i)(v) - \sigma_2(\gamma_i)(v)|.$$

Here $\|\cdot\|_1$ denotes the ℓ_1 -norm on $\mathbb{R}^{|V|}$.

Essentially, the value distance is the cumulative distance between traces at $(\gamma_i)_{i=0}^N$ of the time sequence Γ .

Remark 4. Note that, although the time partitions Γ may vary across traces for BD and RBD (i.e., Γ^{σ_i} may differ from Γ^{σ_j}), they are fixed for VD. This is because allowing variable partitions in VD may reduce meaningful diversity: two identical signals could appear artificially distant if compared at different time instants.

The diversified trace synthesis procedure is iterative: at each step it maximizes the value distance between the candidate trace and all previously generated traces, subject to the constraint that all traces share the same time sequences. Let σ_m denote the piecewise-linear signal produced at iteration m , and let $(\gamma_0^{(m)}, \dots, \gamma_N^{(m)})$ be the corresponding δ -stable partition at iteration m . The iterative synthesis procedure is as follows.

- **Initialization** ($m = 1$). Solve the original MILP problem with objective $f_1 := 0$.

- **Iteration** ($m > 1$). Let \mathcal{C} denote the original MILP constraints from Sect. 3. Solve the following MILP:

$$\begin{aligned}
 & \text{maximize} && \sum_{k=1}^{m-1} \sum_{i=0}^N \sum_{v \in V} Z_{k,i,v}, && (4) \\
 & \text{subject to} && (\mathcal{C}), \\
 & && \gamma_i = \gamma_i^{(m-1)} && 0 \leq i \leq N, \\
 & && Z_{k,i,v} = \mathbf{abs}(Y_{k,i,v}) && 1 \leq k \leq m-1, 0 \leq i \leq N, v \in V, \\
 & && Y_{k,i,v} = x_{i,v} - \sigma_k(\gamma_i)(v) && 1 \leq k \leq m-1, 0 \leq i \leq N, v \in V.
 \end{aligned}$$

Here, $x_{i,v}$ is the MILP decision variable representing the value of signal variable $v \in V$ on the interval $[\gamma_{i-1}, \gamma_i]$.

At iteration $m > 1$, to i) ensure that synthesized traces share the same δ -stable partition, and ii) simplify the encoding of absolute values in the objective, we augment \mathcal{C} with the constraints above. The relations $Z_{k,i,v} = \mathbf{abs}(Y_{k,i,v})$ are implemented by standard MILP linearization (introducing auxiliary variables and linear inequalities), so the resulting problem remains an MILP.

Remark 5. Due to the use of absolute-value terms in the constraints, the MILP complexity increases with each iteration. Specifically, at iteration m , new variables $Y_{m,i,v}$ and $Z_{m,i,v}$ are introduced for every $i = 0, \dots, N$ and $v \in V$, adding $2(N+1)|V|$ variables per iteration. Consequently, the overall complexity of generating m solutions grows faster than linearly with respect to the size of the individual MILP problem.

Nevertheless, the value distance method is substantially faster than the two Boolean distance approaches when synthesizing a small number of traces. The value distance objective sums over signal variables, whereas the Boolean distance objective sums over subformulas. Since $|V|$ is typically much smaller than $|\text{Sub}(\varphi)|$ for realistic STL formulas, the value distance objective contains far fewer terms and is therefore cheaper to optimize. See Sect. 6 for experimental results.

5 Demonstration of STLts-Div

STLts-Div is a Python package for MILP-based diversified trace synthesis from STL. It can be installed by i) ensuring a working Gurobi installation, and ii) running `pip install stltspref`. We use the name `stltspref` for the pip package to accommodate future extensions incorporating preference-based trace synthesis methods.

The procedure to run STLts-Div is shown in Fig. 1. The user must provide the following input.

```

1 import gurobipy as gp
2 from stltspref.problem import create_stl_milp_problem
3
4 bound = 10                                # variability bound (see Problem 1)
5 milp = gp.Model()                          # Gurobi model
6 prob = create_stl_milp_problem(milp, N=bound) # initialize the MILP problem

```

(a) Create the MILP problem `prob`.

```

1 # Define system model and its continuous states
2 car = prob.create_system_model()
3 car.add_state('x', 0, 40)                # position x in [0,40]
4 car.add_state('v', -5, 5)                 # velocity v in [-5,5]
5 car.add_state('a', -3, 3)                 # acceleration in [-3,3]
6
7 # Set initial state and dynamics (double-integrator model)
8 car.set_initial_state(x=(0,0), v=(0,0), a=(0,0)) # initial value ranges
9 car.add_dynamics('a', -3, 3, constant=True)
10 car.add_double_integrator_dynamics('x', 'v', 'a')

```

(b) Define the system model $\dot{x} = v$, $\dot{v} = a$, $a \in [-3, 3]$, and add it to `prob`.

```

1 from stltspref.linear_expression import LinearExpression as L
2 from stltspref.stl import Atomic, BoundedAlw, Ev
3
4 danger = Atomic(L.unit('x') <= 10) # atomic proposition
5 stl_spec = Ev(BoundedAlw([0, 5], danger)) # STL specification
6
7 prob.initialize_milp_formulation(stl_spec) # add the spec to MILP

```

(c) Specify the STL formula $\diamond(\square_{[0,5]} x \leq 10)$ and add it to `prob`.

```

1 from stltspref.preferential_synthesis import diversity_finder
2
3 diversity_method = 'RBD' # random Boolean distance method
4 trace_num = 5           # synthesis 5 diverse traces
5
6 diversity_finder(prob, diversity_method, bound, trace_num)

```

(d) Solve `prob` to synthesize diverse traces.

Fig. 1. Workflow for MILP-based trace synthesis using STLts-Div.

- A system model. STLts-Div supports models such as *double-integrator dynamics* and *rectangular hybrid automata* (RHA) (cf. [20] for details); see Fig. 1a.
- An STL formula (encoded in a DSL style); see Fig. 1c;
- A diversification method: SP (Gurobi solution pool [10]), BD (Boolean distance; Sect. 4.1), RBD (randomized Boolean distance; Sec. 4.2), or VD (value difference; Sect. 4.3); see Fig. 1d;
- The desired number of traces to synthesize (a positive integer); see Fig. 1d.

6 Experimental Analysis

In this section, we evaluate our tool across different benchmarks and STL specifications to assess the efficiency and effectiveness of each diversification method. As a baseline, we also include the standard approach that extracts traces directly from the solution pool of Gurobi.

6.1 Experiment Setting

Benchmarks. To evaluate the efficiency of our diversification tool, we used benchmarks from four categories: *Danger-then-Stop* (DSTOP), *Rear-End Near Collision* (RNC), *Navigation* (NAV), and *Disturbance Scenarios in ISO 34502* (ISO). Each category shares the same system model but includes several distinct STL specifications, resulting in a total of 13 benchmarks.

Danger then stop (DSTOP) Consider two cars in a single lane, where the rear car Car_r repeatedly approaches the front car Car_f too closely, creating a potentially dangerous situation. Eventually, one of the cars comes to a stop. The system dynamics are given by $\dot{x}_r = v_r$, $\dot{v}_r = a_r$, $\dot{x}_f = v_f$, $\dot{v}_f = a_f$, with $a_r, a_f \in [-3, 3]$. The STL specification DSTOP is defined as follows:

$$\begin{aligned} \text{lead} &::= x_f - x_r \geq 0, & \text{close} &::= x_f - x_r \leq 10, & \text{far} &::= x_f - x_r \geq 40, \\ \text{DSTOP} &::= (\Box \text{lead}) \wedge (\Box_{[0,2]} \text{far}) \wedge [(\Diamond_{[0,10]} \text{close}) \mathcal{U}_{[10,15]} (v_f \leq 0 \vee v_r \leq 0)]. \end{aligned}$$

Rear-End Near Collision (RNC1–3) The system model for RNC1–3 is identical to that of DSTOP. These benchmarks describe driving scenarios in which the rear car Car_r approaches the front car Car_f too closely. The corresponding STL specifications, RNC1, RNC2, and RNC3, are defined as follows:

$$\begin{aligned} \text{danger} &::= x_f - x_r \leq 10, \\ \text{dyn_inv} &::= x_f - x_r \geq 0 \wedge 2 \leq v_f \leq 27 \wedge 2 \leq v_r \leq 27, \\ \text{trimming} &::= (\Diamond \text{danger}) \Rightarrow ((\Box_{[0,0.2]} a_r \geq 0.5) \mathcal{U} \text{danger}), \\ \text{trimming2} &::= (\Diamond \text{danger}) \Rightarrow ((\Box_{[0,1]} a_r \geq 1) \mathcal{U} \text{danger}), \\ \text{RNC1} &::= \Box(\text{dyn_inv} \wedge \text{trimming}) \wedge \Diamond_{[0,9]} \Box_{[0,1]} \text{danger}, \\ \text{RNC2} &::= (\Box(x_f - x_r \geq 0)) \wedge \\ &\quad \Diamond_{[0,9]} ((\Box_{[0,1]} \text{danger}) \wedge (\Box_{[0,1]} a_r \geq 1) \wedge (\Diamond_{[1,5]} \neg \text{danger})), \\ \text{RNC3} &::= \Box(\text{dyn_inv} \wedge \text{trimming2}) \wedge \Diamond_{[0,9]} \Box_{[0,1]} \text{danger}. \end{aligned}$$




Fig. 2. The NAV1-2 RHA

Navigation (NAV1–2) The system model for NAV1–2 is adapted from [7], a standard example of a rectangular hybrid automaton (RHA) widely used in prior work (e.g., [4, 20]). The model describes the motion of a robot on a grid consisting of four rectangular regions ℓ_1, \dots, ℓ_4 . Each region has its own dynamics (see Fig. 2). The environment includes an unsafe region `unsafeR` (red) and a goal region `goalR` (green). The robot starts from an initial position (x_0, y_0) where $x_0 \in [0, 3]$ and $y_0 = 0$ (blue line). The STL specifications are $\text{NAV1} ::= \diamond(\square_{[0,3]}((x, y) \in \text{goalR})) \wedge \square(x \notin \text{unsafeR})$ and $\text{NAV2} ::= \square((x, y) \in \ell_3 \Rightarrow \diamond_{[0,3]}(x, y) \in \ell_4)$. Here, NAV1 specifies a reach-avoid property, while NAV2 requires that whenever the robot enters region ℓ_3 , it must reach region ℓ_4 within three seconds.

Disturbance Scenarios in ISO 34502 (ISO1, ISO3, ..., ISO8) ISO 34502 is an international standard for testing autonomous vehicles. The system model is similar to that of RNC1–3, and the specifications `ISO1`, `ISO3`, ..., `ISO8` are obtained in [19]. `ISO2` is omitted because it involves three vehicles. Each spec `ISO i` follows the general structure defined in [19]:

$$\begin{aligned} \text{ISO}_i &::= \text{initSafe} \wedge \text{disturb}_i, \\ \text{disturb}_i &::= \text{initialCondition}_i \wedge \text{behaviourSV}_i \wedge \text{behaviourPOV}_i, \end{aligned}$$

where `SV` denotes the subject (“ego”) vehicle and `POV` the principal other vehicle. The subformulas `initialCondition $_i$` , `behaviourSV $_i$` , and `behaviourPOV $_i$` vary across scenarios (indexed by i). Their detailed definitions are beyond the scope of this paper; please refer to [19] for full descriptions.

Experiment design. We evaluate four diversification strategies implemented in our tool—SP (solution pool), BD (Boolean distance), RBD (randomized Boolean distance), and VD (value distance)—for synthesizing diverse traces. For each benchmark and each method, we generate 8 traces. We set the variability bound as 10 uniformly for each benchmark. While measuring diversity via coverage metrics over the space of all possible traces would provide a comprehensive quantification of diversity, the infinite nature of the trace space (since signals evolve in continuous time) makes such an approach computationally intractable. Investigating tractable approximations is left for future work. Instead, we quantify diversity using the aggregate pairwise Euclidean distance (APED) over time, a common choice in the literature [5]:

$$D_E := \sum_{i < j} \int_0^T \|\sigma_i(t) - \sigma_j(t)\|_2 dt,$$

where i and j index the traces ($i, j \in \{1, \dots, 8\}$ in our experiments) and T is the time horizon. The synthesized traces are resampled over $[0, T]$ with a time step of 0.02, and the integral is computed numerically in Python.

Our experiments were conducted on an Amazon EC2 c5.2xlarge instance (3.00 GHz Intel Xeon Platinum 8275CL, 16 GB RAM) running Ubuntu Server 20.04 LTS.

6.2 Evaluation

We evaluate our tool along from *effectiveness* and *efficiency*.

Effectiveness. Table 1a reports APED for each method and benchmark under a 6-minute per-trace timeout; a dash “—” indicates that no feasible trace was found within the timeout. Across benchmarks, our diversification methods (BD, RBD, VD) consistently achieve substantially higher diversity than the baseline (SP)—typically about an order of magnitude larger APED.

Occasionally the baseline attains $\text{APED} > 0$, reflecting that some specifications are relatively permissive for the given system model: multiple distinct MILP assignments can satisfy φ while respecting the dynamics, yielding non-identical traces. The *DSTOP*, *IS05*, and *IS08* exhibit this behaviour.

Efficiency. Table 1a reports the total time required to generate the full set of traces. For *NAV1* and *NAV2* no feasible trace was found within our setting; this matches [20], which reports that *NAV1* and *NAV2* becomes solvable only with a much larger variability bound (e.g. $N = 17$ for *NAV1*, versus our uniform bound $N=10$).

Because our diversification is posed as an optimization problem, a solver timeout *does not* necessarily imply infeasibility—it may simply return a suboptimal or low-quality solution. To assess sensitivity to the time budget, we repeated the experiments with a reduced per-trace timeout of 60 seconds; results appear in Table 1b.

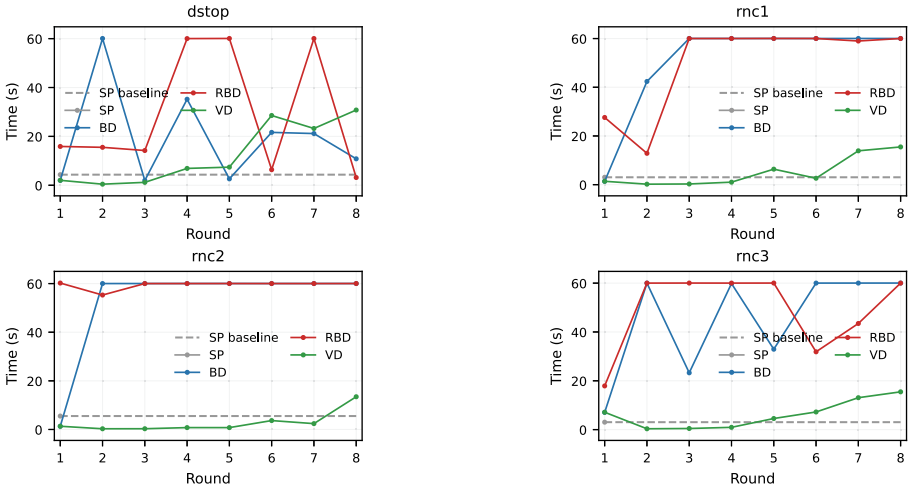


Fig. 3. Round time of *DSTOP*, *RNC1*, *RNC2*, *RNC3*.

For relatively simple benchmarks (e.g., *DSTOP*, *RNC1*, *RNC2*), reducing the timeout has little effect on diversity for any method. In particular, *RBD* requires

Table 1. Aggregated pairwise Euclidean distance under two timeout settings.

(a) Timeout per trace: 360 s.

Benchmarks	SP	BD	RBD	VD
DSTOP	~129k (4.62 s)	~1.6M (200 s)	~934k (1178 s)	~3M (107 s)
RNC1	0(3.11 s)	~1.3M (2207 s)	~1.5M (1881 s)	~1.6M (44.0 s)
RNC2	0(5.87 s)	~2.8M (2523 s)	~2.2M (1867 s)	~4.8M (25.4 s)
RNC3	0(3.13 s)	~1.7M (1473 s)	~1.3M (158 s)	~1.5M (53.0 s)
NAV1	—	—	—	—
NAV2	—	—	—	—
IS01	~18 (95.9 s)	~466k (2529 s)	~403k (2882 s)	~1.5M (2529 s)
IS03	~153 (63.8 s)	~701k (2527 s)	~71k (2882 s)	~792k (2528 s)
IS04	~772 (29.6 s)	~633k (2526 s)	~174k (2881 s)	~598k (2526 s)
IS05	~34k (266 s)	~1M (2561 s)	~32k (2882 s)	~1.7M (2361 s)
IS06	~122 (19.1 s)	~413k (2529 s)	—	~1.6M (1653 s)
IS07	—	~230k (2556 s)	—	~911k (2555 s)
IS08	~1.2k (66.8 s)	~447k (2527 s)	~291k (2882 s)	~896k (2444 s)

(b) Timeout per trace: 60 s.

Benchmarks	SP	BD	RBD	VD
DSTOP	~130k (4.33 s)	~1.6M (157 s)	~1M (237 s)	~3M (102 s)
RNC1	0(3.03 s)	~1.3M (405 s)	~1.4M (401 s)	~1.6M (42 s)
RNC2	0(5.52 s)	~2.4M (423 s)	~2.3M (477 s)	~4.8M (24.4 s)
RNC3	0(3.05 s)	~1.7M (365 s)	~1.2M (395 s)	~1.6M (50.6 s)
NAV1	—	—	—	—
NAV2	—	—	—	—
IS01	—	~79k (429 s)	—	~1.4M (429 s)
IS03	—	0(428 s)	—	~133k (426 s)
IS04	~772 (32.1 s)	~177k (426 s)	—	~256k (425 s)
IS05	—	0(464 s)	—	~1.5M (464 s)
IS06	~122 (21.2 s)	0(429 s)	—	~1.6M (389 s)
IS07	—	0(457 s)	—	~271k (458 s)
IS08	—	~88k (427 s)	0(480 s)	~431k (426 s)

Entries show APED; parentheses give the total time to produce 8 traces (in seconds); ‘—’ indicates that no feasible traces were produced within the timeout; numbers use suffixes $k = 10^3$ and $M = 10^6$;

significantly less time while producing nearly the same APED. This suggests that, with a larger time budget (i.e., 360s per trace), the optimization process may spend most of its time attempting to approach the optimal solution without reaching it before the timeout. In practice, however, a near-optimal solution already provides sufficient diversity, which explains the similar APED values observed under both timeout settings.

By contrast, on more challenging benchmarks (e.g., IS01, IS08, IS04), performance diverges: RBD often fails to produce feasible traces, and BD frequently returns only trivial repeats after the first iteration (APED = 0). VD is the most robust under tight time limits: it continues to produce nontrivial, diverse traces even with the shorter timeout.

To study how per-iteration runtime (i.e., the time to synthesize one trace) evolves with the iteration index, we record the runtime at each iteration for DSTOP, RNC1, RNC2, and RNC3 and plot the results in Fig. 3. For VD, the per-iteration time increases steadily, consistent with our analysis in Sect. 4.3 that each round enlarges the MILP problem (additional auxiliary variables). In contrast, BD and RBD show fluctuating runtimes with no clear trend: BD changes the coefficients of the accumulated objective each round, and RBD re-samples the reference, both of which alter the search landscape and lead to non-monotonic solver times.

Table 2. Comparison of diversification methods

Criteria	SP	BD	RBD	VD
Total runtime	●1-100	○Varying with timeout setting	○Varying with timeout setting	○Varying with timeout setting
Stable per-trace time	-	⊙ Varying across iterations	⊙ Varying across iterations	○ Time increase per iteration
Diversity	○	●	●	●
Compatibility with other tools	● Support all MILP-based formulations	⊙ Compatible only with full Boolean-variable formulations	⊙ Compatible only with full Boolean-variable formulations	○ Limited compatibility
Deterministic	✓	✓	✗	✓

● = good; ⊙ = moderate; ○ = poor; - = not applicable.

Summary. Table 2 contrasts the three diversification strategies. Achieving diversity incurs extra runtime for all of them. The ideas behind BD and RBD are broadly portable: for problems with only Boolean decision variables, adding (possibly randomized) Boolean distance terms can produce multiple diverse solutions. By contrast, VD depends on continuous state variables and absolute-value linearization, which enlarges the MILP problem and makes the approach less directly compatible with other tools or settings.

7 Conclusion

We introduced diversification into MILP-based STL trace synthesis. We formulated several complementary objectives—e.g., Boolean distance, randomized Boolean distance, and value distance—that promote dissimilarity by maximizing an appropriate distance between the candidate trace and previously generated

ones. We implemented these objectives in STLts-Div and evaluated effectiveness and efficiency on a suite of benchmarks, showing clear improvements over a solution-pool baseline.

A complete MILP encoding of robust semantics remains open. Such an encoding could enable richer preference objectives and stronger guarantees, including true corner-case synthesis (robustness exactly 0) and additional forms of preferential synthesis.

References

1. Alur, R., Feder, T., Henzinger, T.A.: The benefits of relaxing punctuality. *Journal of the ACM (JACM)* **43**(1), 116–146 (1996)
2. Atkins, E.M., Bradley, J.M.: Aerospace cyber-physical systems education. In: *AIAA Infotech@ Aerospace (I@ A) Conference*. p. 4809 (2013)
3. Bae, K., Lee, J.: Bounded model checking of signal temporal logic properties using syntactic separation. *Proceedings of the ACM on Programming Languages* **3**(POPL), 1–30 (2019)
4. Bu, L., Frehse, G., Kundu, A., Ray, R., Shi, Y., Zaffanella, E., et al.: Arch-comp22 category report: Hybrid systems with piecewise constant dynamics and bounded model checking. In: *Proceedings of 9th International Workshop on Applied Verification of Continuous and Hybrid Systems (ARCH22)*. EPiC Series in Computing. vol. 90, pp. 44–57 (2022)
5. Cheng, M., Zhou, Y., Xie, X.: Behavexplor: Behavior diversity guided testing for autonomous driving systems. In: *Proceedings of the 32nd ACM SIGSOFT International Symposium on Software Testing and Analysis*. pp. 488–500 (2023)
6. Donzé, A., Raman, V., Frehse, G., Althoff, M.: Blustl: Controller synthesis from signal temporal logic specifications. *ARCH@ CPSWeek* **34**, 160–8 (2015)
7. Duggirala, P.S., Mitra, S.: Abstraction refinement for stability. In: *2011 IEEE/ACM Second International Conference on Cyber-Physical Systems*. pp. 22–31. IEEE (2011)
8. Ernst, G., Arcaini, P., Bennani, I., Chandratre, A., Donzé, A., Fainekos, G., Frehse, G., Gaaloul, K., Inoue, J., Khandait, T., et al.: Arch-comp 2021 category report: Falsification with validation of results. In: *ARCH@ ADHS*. pp. 133–152 (2021)
9. Fainekos, G.E., Pappas, G.J.: Robustness of temporal logic specifications for continuous-time signals. *Theoret. Comput. Sci.* **410**(42), 4262–4291 (Sep2009). [10.1016/j.tcs.2009.06.021](https://doi.org/10.1016/j.tcs.2009.06.021), <http://dx.doi.org/10.1016/j.tcs.2009.06.021>
10. Gurobi Optimization, LLC: Gurobi Optimizer Reference Manual (2024), <https://www.gurobi.com>
11. Jouve-Genty, M., Su, H., Sato, S., An, J., Zhang, Z., Hasuo, I.: STLts-Div: Diversified trace synthesis from STL specifications using MILP (Extended Version) (2025), available on arXiv
12. Lee, J., Yu, G., Bae, K.: Efficient smt-based model checking for signal temporal logic. In: *2021 36th IEEE/ACM international conference on automated software engineering (ASE)*. pp. 343–354. IEEE (2021)
13. Lindemann, L., Dimarogonas, D.V.: Control barrier functions for signal temporal logic tasks. *IEEE control systems letters* **3**(1), 96–101 (2018)
14. Maler, O., Nickovic, D.: *Monitoring Temporal Properties of Continuous Signals*, p. 152–166. Springer Berlin Heidelberg (2004)https://doi.org/10.1007/978-3-540-30206-3_12, http://dx.doi.org/10.1007/978-3-540-30206-3_12

15. Prabhakar, P., Lal, R., Kapinski, J.: Automatic trace generation for signal temporal logic. In: 2018 IEEE Real-Time Systems Symposium (RTSS). pp. 208–217. IEEE (2018)
16. Rabinovich, A.M.: On the decidability of continuous time specification formalisms. *J. Log. Comput.* **8**(5), 669–678 (1998)
17. Raman, V., Donzé, A., Sadigh, D., Murray, R.M., Seshia, S.A.: Reactive synthesis from signal temporal logic specifications. In: Proceedings of the 18th international conference on hybrid systems: Computation and control. pp. 239–248 (2015)
18. Raman, V., Maasoumy, M., Donzé, A.: Model predictive control from signal temporal logic specifications: A case study. In: Proceedings of the 4th ACM SIGBED international workshop on design, modeling, and evaluation of cyber-physical systems. pp. 52–55 (2014)
19. Reimann, J., Mansion, N., Haydon, J., Bray, B., Chattopadhyay, A., Sato, S., Waga, M., André, É., Hasuo, I., Ueda, N., et al.: Temporal logic formalisation of iso 34502 critical scenarios: modular construction with the rss safety distance. In: Proceedings of the 39th ACM/SIGAPP Symposium on Applied Computing. pp. 186–195 (2024)
20. Sato, S., An, J., Zhang, Z., Hasuo, I.: Optimization-Based Model Checking and Trace Synthesis for Complex STL Specifications, p. 282–306. Springer Nature Switzerland (2024) https://doi.org/10.1007/978-3-031-65633-0_13, http://dx.doi.org/10.1007/978-3-031-65633-0_13
21. Su, H., Feng, S., Zhan, S., Zhan, N.: Switching controller synthesis for hybrid systems against stl formulas. In: International Symposium on Formal Methods. pp. 229–247. Springer (2024)
22. Su, H., Shankar, S., Pinisetty, S., Roop, P.S., Zhan, N.: Runtime enforcement of cps against signal temporal logic. In: Proceedings of the 28th ACM International Conference on Hybrid Systems: Computation and Control. pp. 1–11 (2025)
23. Tomlin, C.J., Lygeros, J., Sastry, S.S.: A game theoretic approach to controller design for hybrid systems. *Proc. IEEE* **88**(7), 949–970 (2000)
24. Ye, P., Entcheva, E., Smolka, S.A., Grosu, R.: Modelling excitable cells using cycle-linear hybrid automata. *IET Syst. Biol.* **2**(1), 24–32 (2008)
25. Zhang, Z., An, J., Arcaini, P., Hasuo, I.: Online causation monitoring of signal temporal logic. In: International Conference on Computer Aided Verification. pp. 62–84. Springer (2023)

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

