

Formal Requirement Debugging for Testing and Verification of Cyber-Physical Systems

ADEL DOKHANCHI, Arizona State University

BARDH HOXHA, Southern Illinois University

GEORGIOS FAINEKOS, Arizona State University

A framework for the elicitation and debugging of formal specifications for Cyber-Physical Systems is presented. The elicitation of specifications is handled through a graphical interface. Two debugging algorithms are presented. The first checks for erroneous or incomplete temporal logic specifications without considering the system. The second can be utilized for the analysis of reactive requirements with respect to system test traces. The specification debugging framework is applied on a number of formal specifications collected through a user study. The user study establishes that requirement errors are common and that the debugging framework can resolve many insidious specification errors.

CCS Concepts: • **Theory of computation** → **Logic and verification**; • **Computing methodologies** → **Model verification and validation**;

Additional Key Words and Phrases: MITL, SAT, LTL, STL, SMT, CPS

ACM Reference format:

Adel Dokhanchi, Bardh Hoxha, and Georgios Fainekos. 2017. Formal Requirement Debugging for Testing and Verification of Cyber-Physical Systems. *ACM Trans. Embed. Comput. Syst.* 17, 2, Article 34 (December 2017), 26 pages.

<https://doi.org/10.1145/3147451>

1 INTRODUCTION

Testing and verification of Cyber-Physical Systems (CPSs) is important due to the safety-critical applications of CPSs such as medical devices and transportation systems. It has been shown that utilizing formal specifications can lead to improved testing and verification (Fainekos et al. 2012; Jin et al. 2013; Yang et al. 2012; Kapinski et al. 2015). However, developing formal specifications using logics is a challenging and error-prone task even for experts who have formal mathematical training. Therefore, in practice, system engineers usually define specifications in natural language. Natural language is convenient to use in many stages of system development, but its inherent ambiguity, inaccuracy, and inconsistency make it unsuitable for use in defining specifications.

To assist in the elicitation of formal specifications, in Hoxha et al. (2014b, 2015), we presented a graphical formalism and the corresponding tool ViSPEC that can be utilized by users in both

This work was partially supported by NSF CNS 1116136, NSF CNS 1350420, NSF IIP-1361926, and the NSF I/UCRC Center for Embedded Systems.

Authors' addresses: A. Dokhanchi and G. Fainekos, School of Computing, Informatics and Decision Systems Engineering, Arizona State University, Tempe, AZ 85281 USA; emails: {adokhanc, fainekos}@asu.edu; B. Hoxha, Department of Computer Science, Southern Illinois University, Carbondale, IL 62901 USA; email: bhoxha@cs.siu.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 ACM 1539-9087/2017/12-ART34 \$15.00

<https://doi.org/10.1145/3147451>

academia and industry. Namely, a user-developed graphical input is translated to a Metric Interval Temporal Logic (MITL) formula. The formal specifications in MITL can be used for testing and verification with tools such as S-TALiRo (Annapureddy et al. 2011) and Breach (Donze 2010).

In Hoxha et al. (2015), the tool was evaluated through a usability study, which showed that ViSPEC users were able to use the tool to elicit formal specifications. The usability study results also indicated that in a few cases, the developed specifications were incorrect. This raised two questions. First, are these issues artifacts of the graphical user interface? Second, can we automatically detect and report issues with the requirements themselves?

We have created an online survey¹ to answer the first question. Namely, we conducted a usability study on MITL by targeting users with working knowledge in temporal logics. In our online survey, we tested how well formal method users can translate natural requirements to MITL. That is, given a set of requirements in natural language, users were asked to formalize the requirements in MITL. The study is ongoing, but preliminary results indicate that even users with working knowledge of MITL can make errors in their specifications.

For example, for the natural language specification “At some time in the first 30 seconds, the vehicle speed (v) will go over 100 and stay above 100 for 20 seconds,” the specification $\varphi = \Diamond_{[0,30]}((v > 100) \Rightarrow \Box_{[0,20]}(v > 100))$ was provided as an answer by a user with formal logic background. Here, $\Diamond_{[0,30]}$ stands for “eventually within 30 time units” and $\Box_{[0,20]}$ for “always from 0 to 20 time units.” However, the specification φ is a *tautology*!; i.e., it evaluates to true no matter what the system behavior is and thus the requirement φ is invalid. This is because, if at some time t between 0 and 30 seconds the predicate $(v > 100)$ is false, then the implication (\Rightarrow) will trivially evaluate to true at time t and thus φ will evaluate to true as well. On the other hand, if the predicate $(v > 100)$ is true for all times between 0 and 30 seconds, then the subformula $\Box_{[0,20]}(v > 100)$ will be true at all times between 0 and 10 seconds. This means that the subformula $(v > 100) \Rightarrow \Box_{[0,20]}(v > 100)$ is true at all times between 0 and 10 seconds. Thus, again, φ evaluates to true, which means that φ is a tautology.

This implies that specification issues are not necessarily artifacts of the graphical user interface and that they can happen even for users who are familiar with temporal logics. Hence, specification elicitation can potentially become an issue as formal and semiformal testing and verification methods and tools are being adopted by industry. This is because specification elicitation can be performed by untrained users. Therefore, effort can be wasted in checking incorrect requirements, or even worse, the system can pass the incorrect requirements. Clearly, this can lead to a false sense of system correctness, which leads us to the second question: What can be done in an automated way to prevent specification errors in CPS?

In this work, we have developed a specification debugging framework to assist in the elicitation of formal requirements. The specification debugging algorithm identifies some of the logical issues in the specifications, but not all of them. Namely, it performs the following:

- (1) Validity detection: the specification is unsatisfiable or a tautology.
- (2) Redundancy detection: the formula has redundant conjuncts.
- (3) Vacuity detection: some subformulas do not affect the satisfiability of the formula.

Redundancy and vacuity issues usually indicate some misunderstanding in the requirements. As a result, a wide class of specification errors in the elicitation process can be corrected before any test and verification process is initiated. However, some specification issues cannot be detected unless we consider the system and test the system behaviors with respect to the specification. We provide algorithms to detect specification vacuity with respect to system traces in order to help the

¹The on-line anonymous survey is available through: <http://goo.gl/forms/YW0reiDtgi>.

CPS developer find more vacuity issues during system testing. Our framework can help developers correct their specifications as well as find more subtle errors during testing.

This article is an extended version of the conference paper that appeared in MEMOCODE 2015 (Dokhanchi et al. 2015).

Summary of Contributions:

- (1) We present a specification debugging algorithm for a fragment of MITL (Alur et al. 1996) specifications.
- (2) Using (1), we provide a debugging algorithm for Signal Temporal Logic Specifications (Maler and Nickovic 2004).
- (3) We extend Linear Temporal Logic (LTL) (Clarke et al. 1999) vacuity detection algorithms (Chockler and Strichman 2009) to real-time specifications in MITL.
- (4) We formally define signal vacuity and we provide an algorithm to detect the system traces that vacuously satisfy the real-time specifications in MITL.
- (5) We present experimental results on specifications that typically appear in requirements for CPS.

The above contributions can help us address and solve some of the logical issues that may be encountered when writing MITL specifications. In particular, we believe that our framework will primarily help users with minimal training in formal requirements who use graphical specification formalisms like ViSPEC (Hoxha et al. 2015). The users of ViSPEC can benefit from our feedback and fix any reported issues. In addition, we can detect potential system-level issues using algorithms to determine specification vacuity with respect to system traces during testing.

In this article, the new results over the conference version of the article (Dokhanchi et al. 2015) concern item (4) in the list above and are presented in Sections 6, 7.3, and 7.4 and Appendix B. In addition, we have expanded some examples and added further experimental results that did not appear in Dokhanchi et al. (2015). Furthermore, we added new algorithms in Section 4.

2 RELATED WORKS

The challenge of developing formal specifications has been studied in the past. The most relevant works appear in Autili et al. (2007) and Zhang et al. (2010). Autili et al. (2007) extend Message Sequence Charts and UML 2.0 Interaction Sequence Diagrams to propose a scenario-based formalism called a Property Sequence Chart (PSC). The formalism is mainly developed for specifications on concurrent systems. In Zhang et al. (2010), PSC is extended to Timed PSC, which enables the addition of timing constructs to specifications. Another nongraphical approach to the specification elicitation problem utilizes specification patterns (Dwyer et al. 1998). The patterns provided include commonly used abstractions for different logics including LTL, Computation Tree Logic (CTL), or Quantified Regular Expressions (QREs). This work was extended to the real-time domain (Konrad and Cheng 2005).

Specification debugging can also be considered in areas such as system synthesis (Raman and Kress-Gazit 2011) and software verification (Ammons et al. 2003). In system synthesis, realizability is an important factor, which checks whether the system is implementable given the constraints (environment) and requirements (specification) (Ehlers and Raman 2014; Könighofer et al. 2013; Cimatti et al. 2008; Raman and Kress-Gazit 2011). Specification debugging can also be considered with respect to the environment for robot motion planning. Fainekos (2011) and Kim et al. (2012) considered the problem where the original specification is unsatisfiable with the given environment and robot actions. Then they relax the specification in order to render it satisfiable in the given domain.

One of the most powerful verification methods is model checking (Clarke et al. 1999), where a finite state model of the system is evaluated with respect to a specification. For example, let us consider model checking with respect to the LTL formula $\varphi = \Box(req \Rightarrow \Diamond ack)$, which represents the following Request-Response requirement: “if at any time in the future a *request* happens, then from that moment on an *acknowledge* must eventually happen.” Here, φ can be trivially satisfied in all systems in which a *request* never happens. In other words, if the *request* never happens in the model of the system (let’s say due to a modeling error), our goal for checking the reaction of the system (issuing the *acknowledge*) is not achieved. Thus, the model satisfies the specification but not in the intended way. This may hide actual problems in the model.

Such satisfactions are called *vacuous* satisfactions. Antecedent failure was the first problem that raised vacuity as a serious issue in verification (Beatty and Bryant 1994; Ben-David et al. 2015). Vacuity can be addressed with respect to a model (Beer et al. 2001; Kupferman and Vardi 2003) or without a model (Fisman et al. 2008; Chockler and Strichman 2009). A formula that has a subformula that does not affect the overall satisfaction of the formula is a vacuous formula. It has been proven in Fisman et al. (2008) that a specification φ is satisfied vacuously in all systems that satisfy it if and only if φ is equivalent to some mutations of it. Chockler and Strichman (2009) provide an algorithmic approach to detecting vacuity and redundancy in LTL specifications. Vacuity with respect to testing was considered in Ball and Kupferman (2008). Ball and Kupferman (2008) defined *weak vacuity* for test suites that vacuously pass LTL monitors (e.g., Havelund and Rosu (2004)). The main idea behind the work in Ball and Kupferman (2008) is that some transitions are removed from the LTL specification automata in order to find vacuous passes during testing. Ball and Kupferman (2008) renamed the vacuity in model checking as *strong vacuity*. Post et al. (2011) consider the problem of vacuity detection in the set of requirements formalized in Duration Calculus (Meyer et al. 2008).

Our work extends Chockler and Strichman (2009) and it is applied to a fragment of MITL. We provide a new definition of vacuity with respect to Boolean or real-valued signals. To the best of our knowledge, vacuity of real-time properties such as MITL has not been addressed yet. Although this problem is computationally hard, in practice, the computation problem is manageable, due to the small size of the formulas.

3 PRELIMINARIES

In this work, we take a general approach in modeling CPSs. In the following, \mathbb{R} is the set of real numbers, \mathbb{R}_+ is the set of nonnegative real numbers, \mathbb{Q} is the set of rational numbers, and \mathbb{Q}_+ is the set of nonnegative rational numbers. Given two sets A and B , B^A is the set of all functions from A to B ; i.e., for any $f \in B^A$, we have $f : A \rightarrow B$. We define 2^A to be the power set of set A . Since we primarily deal with bounded time signals, we fix the variable $T \in \mathbb{R}_+$ to denote the maximum time of a signal.

3.1 Metric Interval Temporal Logic

Metric Temporal Logic (MTL) was introduced in Koymans (1990) in order to reason about the quantitative timing properties of Boolean signals. MITL is MTL where the timing constraints are not allowed to be singleton sets (Alur et al. 1996). In the rest of the article, we restrict our focus to a fragment of MITL called Bounded-MITL(\Diamond, \Box), where the only temporal operators allowed are *Eventually* (\Diamond) and *Always* (\Box) operators with timing intervals. Formally, the syntax of Bounded-MITL(\Diamond, \Box) is defined by the following grammar:

Definition 3.1 (Bounded-MITL(\Diamond, \Box) syntax).

$$\phi ::= \top \mid \perp \mid a \mid \neg a \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \Diamond_I \phi_1 \mid \Box_I \phi_1,$$

where $a \in AP$, AP is the set of atomic propositions, \top is true, and \perp is false. Also, I is a nonsingular interval over \mathbb{Q}_+ with defined endpoints. The interval I is right-closed. We interpret MITL semantics over timed traces. A timed trace is a mapping from the bounded real line to sets of atomic propositions ($\mu : [0, T] \rightarrow 2^{AP}$). We assume that the traces satisfy the finite variability condition (non-Zeno condition).²

Definition 3.2 (Bounded-MITL(\diamond, \square) semantics in Negation Normal Form (NNF)). Given a timed trace $\mu : [0, T] \rightarrow 2^{AP}$ and $t, t' \in [0, T]$, and an MITL formula ϕ , the satisfaction relation $(\mu, t) \models \phi$ is inductively defined as:

- $(\mu, t) \models \top$
- $(\mu, t) \models a$ if and only if $a \in \mu(t)$
- $(\mu, t) \models \neg a$ if and only if $a \notin \mu(t)$
- $(\mu, t) \models \phi_1 \wedge \phi_2$ if and only if $(\mu, t) \models \phi_1$ and $(\mu, t) \models \phi_2$
- $(\mu, t) \models \phi_1 \vee \phi_2$ if and only if $(\mu, t) \models \phi_1$ or $(\mu, t) \models \phi_2$
- $(\mu, t) \models \diamond_I \phi_1$ if and only if $\exists t' \in (t + I) \cap [0, T]$ s.t. $(\mu, t') \models \phi_1$.
- $(\mu, t) \models \square_I \phi_1$ if and only if $\forall t' \in (t + I) \cap [0, T]$, $(\mu, t') \models \phi_1$.

Given an interval $I = [l, u]$, $(t + I)$ creates a new interval I' where $I' = [l + t, u + t]$. A timed trace μ satisfies a Bounded-MITL(\diamond, \square) formula ϕ (denoted by $\mu \models \phi$) if and only if $(\mu, 0) \models \phi$. False is defined as $\perp \equiv \neg \top$. In this article, we assume that the Bounded-MITL(\diamond, \square) formula is in NNF,³ where the negation operation is only applied on atomic propositions. NNF is easily obtainable by applying DeMorgan's Law, i.e., $\neg \diamond_I \phi \equiv \square_I \neg \phi$ and $\neg \square_I \phi \equiv \diamond_I \neg \phi$. Any *Implication* operation (\Rightarrow) will be rewritten as $\psi \Rightarrow \phi \equiv \psi' \vee \phi$, where $\psi' \equiv \neg \psi$ and also ψ' is in NNF, and NNF formulas only contain the following Boolean operators of (\wedge, \vee). For simplifying the presentation, when we mention MITL, we mean Bounded-MITL(\diamond, \square). Given MITL formulas ϕ and ψ , ϕ satisfies ψ , denoted by $\phi \models \psi$ if and only if $\forall \mu. \mu \models \phi \Rightarrow \mu \models \psi$. Throughout this article, we use $\phi \in \psi$ to denote that ϕ is a **subformula** of ψ .

3.2 Signal Temporal Logic

The logic and semantics of MITL can be extended to real-valued signals through Signal Temporal Logic (STL) (Maler and Nickovic 2004).

Definition 3.3 (Signal Temporal Logic (Maler and Nickovic 2004)). Let $s : [0, T] \rightarrow \mathbb{R}^m$ be a real-valued signal, and $\Pi = \{\pi_1, \dots, \pi_n\}$ be a collection of predicates or Boolean functions of the form $\pi_i : \mathbb{R}^m \rightarrow \mathbb{B}$, where $\mathbb{B} = \{\top, \perp\}$ is a Boolean value.

For any STL formula Φ_{STL} over predicates Π , we can define a corresponding MITL formula Φ_{MITL} over some atomic propositions AP as follows:

- (1) Define a set AP such that for each $\pi \in \Pi$, there exist some $a_\pi \in AP$.
- (2) For each real-valued signal s we define a μ such that $\forall t. a_\pi \in \mu(t)$ if and only if $\pi(s(t)) = \top$.
- (3) $\forall t. (s, t) \models \Phi_{STL}$ if and only if $(\mu, t) \models \Phi_{MITL}$.

The traces resulting from abstractions through predicates of signals from physical systems satisfy the finite variability assumption. For practical applications, the finite variability assumption is satisfied. Since our article focuses on CPS, with the abuse of terminology, we may use signal to refer to both timed traces and signals.

²The satisfiability tools for MITL that we use in Section 7.1 assume that the traces satisfy the finite variability condition (Bersani et al. 2016).

³We relax this assumption for addressing the Request-Response specifications (see Section 6.1).

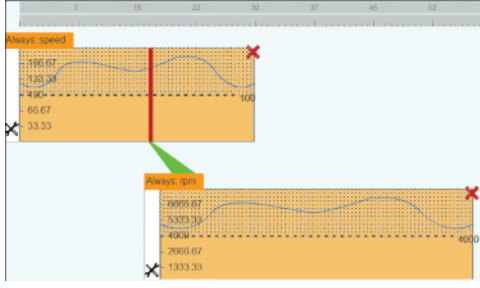


Fig. 1. Graphical representation of $\phi_{cps} = \Box_{[0,30]}((speed > 100) \Rightarrow \Box_{[0,40]}(rpm > 4000))$.

$S \longrightarrow \neg T \mid T$
 $T \longrightarrow A \mid B \mid C$
 $A \longrightarrow P \mid (P \wedge A) \mid (P \Rightarrow A)$
 $B \longrightarrow \Box_I D \mid \Diamond_I D$
 $C \longrightarrow \Box_I \Diamond_I D \mid \Diamond_I \Box_I D$
 $D \longrightarrow p \mid (p \Rightarrow A) \mid (p \wedge A) \mid (p \Rightarrow B) \mid (p \wedge B)$
 $P \longrightarrow p \mid \Box_I p \mid \Diamond_I p$

Fig. 2. ViSPEC grammar to generate MITL.

Table 1. Classes of Specifications Expressible with the Graphical Formalism

Specification Class	Explanation
Safety	Specifications of the form $\Box\phi$ used to define specifications where ϕ should always be true.
Reachability	Specifications of the form $\Diamond\phi$ used to define specifications where ϕ should be true at least once in the future (or now).
Stabilization	Specifications of the form $\Diamond\Box\phi$ used to define specifications that, at least once, ϕ should be true and from that point on stay true.
Oscillation	Specifications of the form $\Box\Diamond\phi$ used to define specifications that, it is always the case, at some point in the future, ϕ repeatedly will become true.
Implication	Specifications of the form $\phi \Rightarrow \psi$ require that ψ should hold when ϕ is true.
Request-Response	Specifications of the form $\Box(\phi \Rightarrow M\psi)$, where M is temporal operator, used to define an implicative response between two specifications where the timing of M is relative to the timing of \Box .
Conjunction	Specifications of the form $\phi \wedge \psi$ used to define the conjunction of two subspecifications.
Nonstrict Sequencing	Specifications of the form $N(\phi \wedge M\psi)$, where N and M are temporal operators, used to define a conjunction between two specifications where the timing of M is relative to the timing of N .

3.3 Visual Specification Tool

The Visual Specification Tool (ViSPEC) (Hoxha et al. 2015) enables the development of formal specifications for CPS. The graphical formalism enables reasoning on both timing and event sequence occurrence. Consider the specification $\phi_{cps} = \Box_{[0,30]}((speed > 100) \Rightarrow \Box_{[0,40]}(rpm > 4000))$. It states that whenever, within the first 30 seconds, *vehicle speed* goes over 100, then from that moment on, the *engine speed (rpm)*, for the next 40 seconds, should always be above 4,000. Here, both the sequence and timing of the events are of critical importance. See Figure 1 for the visual representation of ϕ_{cps} .

Users develop specifications using a visual formalism that can be translated to an MITL formula. The set of specifications that can be generated from this graphical formalism is a proper subset of the set of MITL specifications. Figure 2 represents the grammar that produces the set of formulas that can be expressed by the ViSPEC graphical formalism. In Figure 2, p is an atomic proposition. In the tool, the atomic propositions are automatically derived from graphical templates. For example, the formula $\Box_I \Diamond_I p$ can be generated using the following parse tree $S \longrightarrow T \longrightarrow C \longrightarrow \Box_I \Diamond_I D \longrightarrow \Box_I \Diamond_I p$. ViSPEC provides a variety of templates and the connections between them, which allow the users to express a wide collection of specifications as presented in Table 1. For a more detailed description of ViSPEC, refer to Hoxha et al. (2015).





Fig. 3. Specification elicitation framework.

Fig. 4. The real-valued *Speed* signal and its three Boolean abstractions: $a \equiv \text{speed} > 100$ (solid black line), $b \equiv \text{speed} > 80$ (dotted line), and $c \equiv 100 \geq \text{Speed} > 80$ (gray line).

4 MITL ELICITATION FRAMEWORK

Our framework for elicitation of MITL specifications is presented in Figure 3. Once a specification is developed using ViSPEC, it is translated to STL. Then, we create the corresponding MITL formula from STL. Next, the MITL specification is analyzed by the debugging algorithm, which returns an alert to the user if the specification has inconsistency or correctness issues. The debugging process is explained in detail in the next section.

To enable the debugging of specifications, we must first project the STL predicate expressions (functions) into atomic propositions with independent truth valuations. This is very important because the atomic propositions ($a \in AP$) in MITL are assumed to be independent of each other. However, when we project predicates to the atomic propositions, the dependency between the predicates restricts the possible combinations of truth valuations of the atomic propositions. This notion of predicate dependency is illustrated using the following example. Consider the real-valued signal *Speed* in Figure 4. The Boolean abstraction a (b , respectively) over the *Speed* signal is true when the *Speed* is above 100 (80, respectively). The predicates a and b are related to each other because it is always the case that if $\text{Speed} > 100$, then also $\text{Speed} > 80$. In Figure 4, the Boolean signals for predicates a and b are represented in black solid and dotted lines, respectively. It can be seen that solid and dotted lines are overlapping, which shows the dependency between them. However, this dependency is not captured if we naively substitute each predicate with a unique atomic proposition. If we lose information about the intrinsic logical dependency between a and b , then the debugging algorithm will not find possible specification issues.

For analysis of STL formulas within our MITL debugging process, we must replace the original predicate with nonoverlapping (mutual exclusive) predicates. For the example illustrated in Figure 4, we create a new atomic proposition c that corresponds to $100 \geq \text{speed} > 80$ and the corresponding Boolean signal is represented in gray. In addition, we replace the atomic proposition b with the propositional formula $a \vee c$ since $\text{speed} > 80 \equiv (\text{speed} > 100 \vee 100 \geq \text{speed} > 80)$. Now, the dependency between $\text{Speed} > 100$ and $\text{Speed} > 80$ can be preserved because it is always the case that if a ($\text{Speed} > 100$), then $a \vee c$ ($\text{Speed} > 80$). It can be seen in Figure 4 that the signal b (dotted line) is the disjunction of the solid black (a) and gray (c) signals, where a and c cannot be simultaneously true.

The projection of STL to MITL with independent atomic propositions is conducted using a brute-force approach that runs through all the combinations of predicate expressions to find overlapping parts. The high-level overview of Algorithm 1 is as follows: given the set of predicates $\Pi = \{\pi_1, \dots, \pi_n\}$, the algorithm iteratively calls Algorithm 2 (DecPred) in order to identify predicates whose corresponding sets have nonempty intersections. For each predicate π_i , we assume there exists a corresponding set \mathcal{S}_i such that $\mathcal{S}_i = \{x \mid x \in \mathbb{R}^m, \pi_i(x) = \top\}$. The set \mathcal{S}_i represents part of space \mathbb{R}^m where predicate function π_i evaluates to \top . When no nonempty intersection is found, the Algorithm 1 terminates.

Algorithm 1 creates a temporary copy of Π in a new set Δ . Then in a while-loop Algorithm 2 is called in Line 3 to find overlapping predicates. Algorithm DecPred checks all the combinations of predicates in Δ until it finds overlapping sets (see Line 3). The DecPred partitions two overlapping predicates π_i, π_j into three mutually exclusive predicates $\bar{\pi}_{ij1}, \bar{\pi}_{ij2}, \bar{\pi}_{ij3}$ in Lines 4 to 6. Then π_i, π_j are removed from Δ in Line 7 and new predicates $\bar{\pi}_{ij1}, \bar{\pi}_{ij2}, \bar{\pi}_{ij3}$ are appended to Δ (Line 8). If no overlapping predicates are found, then DecPred returns \emptyset in Line 13 and termCond gets value 1 in Line 7 of Algorithm 1. Now the while-loop will terminate and Ψ contains all nonoverlapping predicates. We must rewrite the predicates in Π with a disjunction operation on the new predicates in Ψ . This operation takes place in Line 11 of Algorithm 1. Since the function CreateDisjunction is trivial, we omit its pseudo-code. The runtime overhead of Algorithm 1 and the size of the resulting set Ψ can be exponential to $|\Pi| = n$, because we can have 2^n possible combinations of predicate evaluations.

ALGORITHM 1: Generate Mutually Exclusive Predicates

Input: Set of predicates $\Pi = \{\pi_1, \dots, \pi_n\}$

Output: Mutually exclusive predicates Ψ

Update Π with Disjunction of predicates Ψ

```

1: termCond  $\leftarrow$  0;  $\Delta \leftarrow \Pi$ 
2: while termCond = 0 do
3:    $\Psi \leftarrow \text{DecPred}(\Delta)$ 
4:   if  $\Psi \neq \emptyset$  then
5:      $\Delta \leftarrow \Psi$ 
6:   else
7:     termCond  $\leftarrow$  1
8:      $\Psi \leftarrow \Delta$ 
9:   end if
10: end while
11:  $\Pi \leftarrow \text{CreateDisjunction}(\Pi, \Psi)$ 
12: return  $\Pi, \Psi$ 
```

ALGORITHM 2: DecPred: Decompose Two Predicates

Input: Set of predicates $\Pi = \{\pi_1, \dots, \pi_n\}$

Output: Set of updated predicates Π

```

1: for  $i = 1$  to size of  $\Pi$  do
2:   for  $j = i + 1$  to size of  $\Pi$  do
3:     if  $\mathcal{S}_{\pi_i} \cap \mathcal{S}_{\pi_j} \neq \emptyset$  then
4:        $\bar{\pi}_{ij1} \leftarrow \mathcal{S}_{\pi_i} \cap \mathcal{S}_{\pi_j}$ 
5:        $\bar{\pi}_{ij2} \leftarrow \mathcal{S}_{\pi_i} \setminus \mathcal{S}_{\pi_j}$ 
6:        $\bar{\pi}_{ij3} \leftarrow \mathcal{S}_{\pi_j} \setminus \mathcal{S}_{\pi_i}$ 
7:       Remove( $\Pi, \{\pi_i, \pi_j\}$ )
8:       Append( $\Pi, \{\bar{\pi}_{ij1}, \bar{\pi}_{ij2}, \bar{\pi}_{ij3}\}$ )
9:     return  $\Pi$ 
10:   end if
11: end for
12: end for
13: return  $\emptyset$ 
```




Fig. 5. Specification debugging.

5 MITL SPECIFICATION DEBUGGING

In the following, we present algorithms that can detect inconsistency and correctness issues in specifications. This will help the user in the elicitation of correct specifications. Our specification debugging process conducts the following checks in this order: (1) validity, (2) redundancy, and (3) vacuity. In brief, validity checking determines whether the specification is unsatisfiable or a tautology. Namely, if the specification is unsatisfiable, no system can satisfy it, and if it is a tautology, every system can trivially satisfy it. For example, $p \vee \neg p$ is a tautology. If an MITL formula passes the validity checking, this means that the MITL is satisfiable but not a tautology.

Redundancy checking determines whether the specification has any redundant conjunct. For example, in the specification $p \wedge \Box_{[0,10]}p$, the first conjunct is redundant. Sometimes redundancy is related to incomplete or erroneous requirements where the user may have wanted to specify something else. Therefore, the user should be notified. Vacuity checking determines whether the specification has a subformula that does not affect the satisfaction of the specification. For example, $\varphi = p \vee \Diamond_{[0,10]}p$ is vacuous since the first occurrence of p does not affect on the satisfaction of φ . This is a logical issue because a part of the specification is overshadowed by the other components.

The debugging process is presented in Figure 5. The feedback (Revision Necessary) to the user is a textual description about the detail of each issue. First, given a specification, a validity check is conducted. If a formula does not pass the validity check, then it means that there is a major problem in the specification and the formula is returned for revision. Therefore, redundancy and vacuity checks are not relevant at that point and the user is notified that the specification either is unsatisfiable or is a tautology. Similarly, if the specification is redundant, it means that it has a conjunct that does not have any effect on the satisfaction of the specification and we return the redundant conjunct to the user for revision. Lastly, if the specification is vacuous, it is returned with the issue for revision by the user. When vacuity is detected, we return to the user the simplified formula, which is equivalent to the original MITL.

5.1 Redundancy Checking

Recall that a specification has a redundancy issue if one of its conjuncts can be removed without affecting the models of the specification. Before we formally present what redundant requirements are, we have to introduce some notation. We consider specification Φ as a conjunction of MITL subformulas (φ_j):

$$\Phi = \bigwedge_{j=1}^k \varphi_j. \quad (1)$$

To simplify discussion, we will abuse notation and we will associate a conjunctive formula with the set of its conjuncts. That is:

$$\Phi = \{\varphi_j \mid j = 1, \dots, k\} \equiv \varphi_1 \cup \varphi_2 \cup \dots \cup \varphi_k. \quad (2)$$

Similarly, $\{\Phi \setminus \varphi_i\}$ represents the specification Φ where the conjunct φ_i is removed:

$$\{\Phi \setminus \varphi_i\} = \{\varphi_j \mid j = 1, \dots, i-1, i+1, \dots, k\} = \bigwedge_{j=1}^{i-1} \varphi_j \wedge \bigwedge_{j=i+1}^k \varphi_j. \quad (3)$$

ALGORITHM 3: Redundancy Checking

Input: Φ (MITL Specification)

Output: R_φ (redundant conjuncts w.r.t. conjunctions)

```

1:  $R_\varphi \leftarrow \emptyset$ 
2: for each conjunctive subformula  $\phi_i \in \Phi$  do
3:   for each conjunct  $\psi_j \in \phi_i$  do
4:     if  $\{\phi_i \setminus \psi_j\} \models \psi_j$  then
5:        $R_\varphi \leftarrow R_\varphi \cup (\psi_j, \phi_i)$ 
6:     end if
7:   end for
8: end for
9: return  $R_\varphi$ 

```

ALGORITHM 4: Vacuity Checking

Input: Φ (MITL Specification)

Output: V_φ (vacuous formulas)

```

1:  $V_\varphi \leftarrow \emptyset$ 
2: for each formula  $\varphi_i \in \Phi$  do
3:   for each  $l \in \text{litOccur}(\varphi_i)$  do
4:     if  $\Phi \models \varphi_i[l \leftarrow \perp]$  then
5:        $V_\varphi \leftarrow V_\varphi \cup \{\Phi \setminus \varphi_i\} \wedge \varphi_i[l \leftarrow \perp]$ 
6:     end if
7:   end for
8: end for
9: return  $V_\varphi$ 

```

Therefore, $\{\Phi \setminus \varphi_i\}$ represents a conjunctive formula. Redundancy in specifications can appear in practice due to the incremental approach that system engineers take in the development of specifications. Redundancy should be avoided in formal specifications because it increases the overhead of the testing and verification processes. In addition, redundancy can be the result of incorrect translation from natural language requirements. In the following, we consider the redundancy removal algorithm provided in Chockler and Strichman (2009) for LTL formulas and we extend it to support MITL formulas.

Definition 5.1 (Redundancy of Specification). A conjunct φ_i is redundant with respect to Φ if

$$\bigwedge_{\psi \in \{\Phi \setminus \varphi_i\}} \psi \models \varphi_i.$$

To reformulate, φ_i is redundant with respect to Φ if $\{\Phi \setminus \varphi_i\} \models \varphi_i$. For example, in $\Phi = \Diamond_{[0,10]}(p \wedge q) \wedge \Diamond_{[0,10]}p \wedge \Box_{[0,10]}q$, the conjunct $\Diamond_{[0,10]}(p \wedge q)$ is redundant with respect to $\Diamond_{[0,10]}p \wedge \Box_{[0,10]}q$ since $\Diamond_{[0,10]}p \wedge \Box_{[0,10]}q \models \Diamond_{[0,10]}(p \wedge q)$. In addition, $\Diamond_{[0,10]}p$ is redundant with respect to $\Diamond_{[0,10]}(p \wedge q) \wedge \Box_{[0,10]}q$ since $\Diamond_{[0,10]}(p \wedge q) \wedge \Box_{[0,10]}q \models \Diamond_{[0,10]}p$. This method can detect both issues and report them to the user. Algorithm 3 finds redundant conjuncts in the conjunction operation of the following levels:

- (1) Conjunction as the root formula (top level)
- (2) Conjunction in the nested subformulas (lower levels)

In the top level, it provides the list of subformulas that are redundant with respect to the original MITL Φ . In the lower levels, if a specification has nested conjunctive subformulas ($\phi_i \in \Phi$), it will return the conjunctive subformula ϕ_i as well as its redundant conjunct $\psi_j \in \phi_i$. For example, if $\Phi = \Diamond_{[0,10]}(p \wedge \Box_{[0,10]}p)$ is checked by Algorithm 3, then it will return the pair of $(p, p \wedge \Box_{[0,10]}p)$ to represent that p is redundant in $p \wedge \Box_{[0,10]}p$. In Line 5, the pair of (ψ_j, ϕ_i) is interpreted as follows: ψ_j is redundant in ϕ_i .

5.2 Specification Vacuity Checking

Vacuity detection is used to ensure that all the subformulas of the specification contribute to the satisfaction of the specification. In other words, vacuity checking enables the detection

of irrelevant subformulas in the specifications (Chockler and Strichman 2009). For example, consider the STL specification $\phi_{stl} = \Diamond_{[0,10]}((speed > 100) \vee \Diamond_{[0,10]}(speed > 80))$. In this case, the subformula $(speed > 100)$ does not affect the satisfaction of the specification. This indicates that ϕ_{stl} is a vacuous specification. We need to create correct atomic propositions for the predicate expressions of ϕ_{stl} to be able to detect such vacuity issues in MITL formulas. If we naively replace the predicate expressions $speed > 100$ and $speed > 80$ with the atomic propositions a and b , respectively, then the resulting MITL formula will be $\phi_{mitl} = \Diamond_{[0,10]}(a \vee \Diamond_{[0,10]}b)$. However, ϕ_{mitl} is not vacuous. Therefore, we must extract nonoverlapping predicates as explained in Section 4. The new specification $\phi'_{mitl} = \Diamond_{[0,10]}(a \vee \Diamond_{[0,10]}(a \vee c))$, where a corresponds to $speed > 100$ and c corresponds to $100 \geq speed > 80$, is the correct MITL formula corresponding to ϕ_{stl} , and it is vacuous. In the following, we provide the definition of MITL vacuity with respect to a signal:

Definition 5.2 (MITL Vacuity with Respect to Timed Trace). Given a timed trace μ and an MITL formula φ , a subformula ψ of φ does not affect the satisfiability of φ with respect to μ if and only if ψ can be replaced with any subformula θ without changing the satisfiability of φ on μ . A specification φ is satisfied vacuously by μ , denoted by $\mu \models_V \varphi$, if there exists a subformula ψ that does not affect the satisfiability of φ on μ .

In the following, we extend the framework presented in Chockler and Strichman (2009) to support MITL specifications. Let φ be a formula in NNF where only predicates can be in the negated form. A *literal* is defined as a predicate or its negation. For a formula φ , the set of literals of φ is denoted by $literal(\varphi)$ and contains all the literals appearing in φ . For example, if $\varphi = (\neg p \wedge q) \vee \Diamond_{[0,10]}p \vee \Box_{[0,10]}q$, then $literal(\varphi) = \{\neg p, q, p\}$. Literal occurrences, denoted by $litOccur(\varphi)$, are a multiset of literals appearing in some order in φ , e.g., by traversal of the parse tree. For the given example, $litOccur(\varphi) = \{\neg p, q, p, q\}$. For each $l \in litOccur(\varphi)$, we create the mutation of φ by substituting the occurrence of l with \perp . We denote the mutated formula as $\varphi[l \leftarrow \perp]$.

Definition 5.3 (MITL Vacuity w.r.t. Literal Occurrence). Given a timed trace μ and an MITL formula φ in NNF, specification φ is vacuously satisfied by μ if there exists a literal occurrence $l \in litOccur(\varphi)$ such that μ satisfies the mutated formula $\varphi[l \leftarrow \perp]$. Formally, $\mu \models_V \varphi$ if $\exists l \in litOccur(\varphi)$ s.t. $\mu \models \varphi[l \leftarrow \perp]$.

THEOREM 5.4 (MITL INHERENT VACUITY). Assume that the specification Φ is a conjunction of MITL formulas. If $\exists \varphi_i \in \Phi$ and $\exists l \in litOccur(\varphi_i)$, such that $\Phi \models \varphi_i[l \leftarrow \perp]$, then Φ is inherently vacuous.

A specification Φ is *inherently vacuous* if it is equivalent to its simplified mutation, which means that Φ is vacuous independent of any signal or system. Inherent vacuity of LTL formulas is addressed in Fisman et al. (2008) and Chockler and Strichman (2009). The proof of THEOREM 5.4 is a straightforward modification of the proofs given in Chockler and Strichman (2009) and Kupferman and Vardi (2003). For completeness in the presentation, we provide the proof in Appendix A. When we do not have a root-level conjunction in the specification ($\Phi = \varphi^4$), we check the vacuity of the formula with respect to itself. In other words, we check whether the specification satisfies its mutation ($\varphi \models \varphi[l \leftarrow \perp]$). Technically, Algorithm 4 as presented returns a list of all the mutated formulas that are equivalent to the original MITL.

6 SIGNAL VACUITY CHECKING

In the previous section, we addressed specification vacuity without considering the system. However, in many cases, specification vacuity depends on the system. For example, consider the

⁴In this case, we assume $\{\Phi \setminus \varphi\} \equiv \top$ in Line 5 of the Algorithm 4.

Table 2. Task List with Automotive System Specifications Presented in Natural Language

Task	Natural Language Specification
1. Safety	In the first 40 seconds, vehicle speed should always be less than 160.
2. Reachability	In the first 30 seconds, vehicle speed should go over 120.
3. Stabilization	At some point in time in the first 30 seconds, vehicle speed will go over 100 and stay above for 20 seconds.
4. Oscillation	At every point in time in the first 40 seconds, vehicle speed will go over 100 in the next 10 seconds.
5. Oscillation	It is not the case that, for up to 40 seconds, the vehicle speed will go over 100 in every 10-second period.
6. Implication	If, within 40 seconds, vehicle speed is above 100, then within 30 seconds from time 0, engine speed should be over 3,000.
7. Request-Response	If, at some point in time in the first 40 seconds, vehicle speed goes over 80, then from that point on, for the next 30 seconds, engine speed should be over 4,000.
8. Conjunction	In the first 40 seconds, vehicle speed should be less than 100 and engine speed should be under 4,000.
9. Nonstrict sequencing	At some point in time in the first 40 seconds, vehicle speed should go over 80 and then from that point on, for the next 30 seconds, engine speed should be over 4,000.
10. Long sequence	If, at some point in time in the first 40 seconds, vehicle speed goes over 80, then from that point on, if within the next 20 seconds the engine speed goes over 4,000, then, for the next 30 seconds, the vehicle speed should be over 100.

LTL specification $\varphi = \Box(req \Rightarrow \Diamond ack)$. The specification φ does not have an inherent vacuity issue (Fisman et al. 2008). However, if req never happens in any of the behaviors of the system, then the specification φ is vacuously satisfied on this specific system. As a result, it has been argued that it is important to add vacuity detection in the model checking process (Beer et al. 2001; Kupferman and Vardi 2003). We encounter the same issue when we test signals and systems with respect to Request-Response STL/MITL specifications.

6.1 Vacuous Signals

Consider the MITL specification $\varphi = \Box_{[0,5]}(req \Rightarrow \Diamond_{[0,10]}ack)$. This formula will pass the MITL Specification Debugging method presented in Section 5. However, any timed trace μ that does not satisfy req at any point in time during the test will vacuously satisfy φ . We refer to timed traces that do not satisfy the antecedent (precondition) of the subformula as *vacuous timed traces*. Similarly, these issues follow for STL formulas over signals as well. Consider Task 6 in Table 2 with the specification $\psi = \Diamond_{[0,40]}(speed > 100) \Rightarrow \Box_{[0,30]}(rpm > 3000)$. Any real-valued signal s that does not satisfy $\Diamond_{[0,40]}(speed > 100)$ will vacuously satisfy ψ . Finding such signals is important in testing and monitoring, since if a signal s does not satisfy the precondition of an STL/MITL specification, then there is no point in considering s as a useful test.

Definition 6.1 (Vacuous Timed Trace (Signal)). Given an MITL (STL) formula φ , a timed trace μ (signal s) is vacuous if it satisfies the Antecedent Failure mutation of φ .

Antecedent Failure is one of the main sources of vacuity. Antecedent Failure occurs in a Request-Response specification such as $\varphi_{RR} = \Box_{[0,5]}(req \Rightarrow \Diamond_{[0,10]}ack)$. We provide a formula mutation that can detect signal vacuity in Request-Response specifications.

Definition 6.2 (Request-Response MITL). A Request-Response MITL formula φ_{RR} is an MITL formula that has one or more implication (\Rightarrow) operations in positive polarity (without any negation). In addition, for each implication operation, the consequent must have a temporal operator at the top level.

In the Request-Response (RR) specifications (Horn et al. 2015), we define sequential events in a specific order (by using the implication operator). Many practical specification patterns based on the Request-Response format are provided for system properties (Dwyer et al. 1998; Konrad and Cheng 2005). Therefore, we can define a chain of events that the system must respond/react to. In an RR specification such as $\varphi_{RR} = \Box_{[0,5]}(req \Rightarrow \Diamond_{[0,10]}ack)$, the temporal operator for the consequent $\Diamond_{[0,10]}ack$ is necessary, unless the system does not have any time to acknowledge the req . For any trace μ in which req never happens, we can substitute ack by any formula and the specification is still satisfied by μ . Therefore, if the antecedent is failed by a trace μ , then φ_{RR} is vacuously satisfied by μ . For each implication subformula ($\varphi \Rightarrow \psi$), the left operand (φ) is the precondition (antecedent) of the implication. An antecedent failure mutation is a new formula that is created with the assertion that the precondition (φ) never happens. Note that RR specifications should not be translated into NNF. For each precondition φ , we create an antecedent failure mutation $\Box_{I_\varphi}(\neg\varphi)$, where I_φ is called the *effective interval* of φ .

Definition 6.3 (Effective Interval). The effective interval of a subformula is the time interval when the subformula can have an impact on the truth value of the whole MITL (STL) specification.

Each subformula is evaluated only in the time window that is provided by the *effective interval*. For example, for the MITL specification $\varphi \wedge \psi$, the effective interval for both φ and ψ is $[0,0]$, because φ and ψ can change the value of $\varphi \wedge \psi$ only within the interval $[0,0]$. Similarly, for the MITL specification $\Box_{[0,10]}\varphi$, the effective interval of φ is $[0,10]$, since the truth value of φ is observed in the time window of $[0,10]$ for evaluating $\Box_{[0,10]}\varphi$. The effective interval is important for the creation of an accurate antecedent failure mutation. This is because the antecedent can affect the truth value of the MITL formula only if it is evaluated in the effective interval. The effective interval is like a time window to make the antecedent observable for an outside observer the way it is observed by the MITL specification.

ALGORITHM 5: Effective Interval Update $EIU(\varphi, I)$

Input: φ (Parse Tree of the MITL formula), I (Effective Interval)

Output: φ (Updated formula with subformulas annotated with effective intervals)

```

1:  $\varphi.EI \leftarrow I$ 
2: if  $\varphi \equiv \neg\varphi_m$  then
3:    $EIU(\varphi_m, I)$ 
4: else if  $\varphi \equiv \varphi_m \vee \varphi_n$  OR  $\varphi \equiv \varphi_m \wedge \varphi_n$  OR  $\varphi \equiv \varphi_m \Rightarrow \varphi_n$  then
5:    $EIU(\varphi_m, I)$ 
6:    $EIU(\varphi_n, I)$ 
7: else if  $\varphi \equiv \Box_{I'}\varphi_m$  OR  $\varphi \equiv \Diamond_{I'}\varphi_m$  then
8:    $I'' \leftarrow I' \oplus I$ 
9:    $EIU(\varphi_m, I'')$ 
10: end if
11: return  $\varphi$ 

```

The effective interval of MITL formulas can be computed recursively using Algorithm 5. To run Algorithm 5, we must process the MITL formula parse tree.⁵ The algorithm must be initialized with the interval of $[0,0]$ for the top node of the MITL formula, namely, $EIU(\varphi, [0,0])$. This is because,

⁵We assume that the MITL specification is saved in a binary tree data structure where each node is a formula with the left/right child as the left/right corresponding subformula of φ . In addition, we assume that the nodes of φ 's tree contain a field called *EI*, where we annotate the effective interval of φ in *EI*, namely, $\varphi.EI \leftarrow I_\varphi$.

according to the semantics of MITL, the value of the whole MITL formula is only important at time zero. In Line 8 of Algorithm 5, the operator \oplus is used to add two intervals as follows:

Definition 6.4 (\oplus). Given intervals $I = [l, u]$ and $I' = [l', u']$, we define $I'' \leftarrow I \oplus I'$ where $I'' = [l'', u'']$ such that $l'' = l + l'$ and $u'' = u + u'$.

If either I or I' is left open (right open, respectively), then I'' will be left open (right open, respectively).⁶ In Line 1 of Algorithm 5, the input interval I is assigned to the effective interval of φ , namely, $\varphi.EI$. If the top operation of φ is a propositional operation ($\neg, \vee, \wedge, \Rightarrow$), then the I will be propagated to subformulas of φ (see Lines 2–6). If the top operation of φ is a temporal operator ($\Box_{I'}, \Diamond_{I'}$), then the effective interval is modified according to Definition 6.4 and the interval $I'' \leftarrow I \oplus I'$ is propagated to the subformulas of φ . For example, assume that the MITL specification is $\varphi_{RR} = \Box_{[1,2]}(\Diamond_{[3,5]}b \Rightarrow (\Box_{[4,6]}(c \Rightarrow \Diamond_{[0,2]}d)))$. The specification φ has two antecedents, $\alpha_1 = \Diamond_{[3,5]}b$ and $\alpha_2 = c$. The effective intervals of α_1 and α_2 are $I_{\alpha_1} = [0, 0] \oplus [1, 2] = [1, 2]$ and $I_{\alpha_2} = [0, 0] \oplus [1, 2] \oplus [4, 6] = [5, 8]$, respectively. As a result, the antecedent failure mutations are $\Box_{[1,2]}(\neg \Diamond_{[3,5]}b)$ and $\Box_{[5,8]}(\neg c)$, respectively. Algorithm 6 returns the list of antecedent failures AF_φ , namely, all the vacuously satisfied implication subformulas by s . If the AF_φ list is empty, then the signal s is not vacuous. To check whether the signal s satisfies φ 's mutations in Algorithm 6 (Line 5), we should use an offline monitor such as Fainekos et al. (2012).

ALGORITHM 6: Antecedent Failure	ALGORITHM 7: Literal Occurrence Removal
Input: φ_{RR}, s (RR Specification, Signal) Output: AF_φ a list of failed antecedents 1: $AF_\varphi \leftarrow \emptyset$ 2: $EIU(\varphi, [0, 0])$ 3: for each implication $(\varphi_i \Rightarrow \psi_i) \in \varphi_{RR}$ do 4: $I\varphi_i \leftarrow \varphi_i.EI$ 5: if $s \models \Box_{I\varphi_i}(\neg \varphi_i)$ then 6: $AF_\varphi \leftarrow AF_\varphi \cup (\varphi_i \Rightarrow \psi_i)$ 7: end if 8: end for 9: return AF_φ	Input: Φ, s (Specification, Signal) Output: MF_φ a list of mutated formulas 1: $MF_\varphi \leftarrow \emptyset$ 2: for each formula $\varphi_i \in \Phi$ do 3: for each $l \in litOccur(\varphi_i)$ do 4: if $s \models \varphi_i[l \leftarrow \perp]$ then 5: $MF_\varphi \leftarrow MF_\varphi \cup \{\varphi_i[l \leftarrow \perp]\}$ 6: end if 7: end for 8: end for 9: return MF_φ

6.2 Vacuity Detection in Testing and Falsification

Detecting vacuous satisfaction of specifications is usually applied on top of model checking tools for finite-state systems (Beer et al. 2001; Kupferman and Vardi 2003). However, in general, the verification problem for hybrid automata (a mathematical model of CPS) is undecidable (Alur et al. 1995). Therefore, a formal guarantee about the correctness of CPS modeling and design is impossible, in general. CPSs are usually safety-critical systems and the verification and validation of these systems is necessary. One approach is to use Model-Based Design (MBD) with a mathematical model of the CPS to facilitate the system analysis and implementation (Abbas et al. 2013). Thus, semiformal verification methods are gaining popularity (Kapinski et al. 2015). Although we cannot solve the correctness problem with testing and monitoring, we can detect possible errors with respect to STL requirements.

In Figure 6, a testing approach for signal vacuity detection is presented. The input generator creates initial conditions and inputs to the system under test. The system under test can be a Model,

⁶Although we assume in Definition 3.1 that intervals are right closed, Algorithm 5 can be applied to right-open intervals as well.

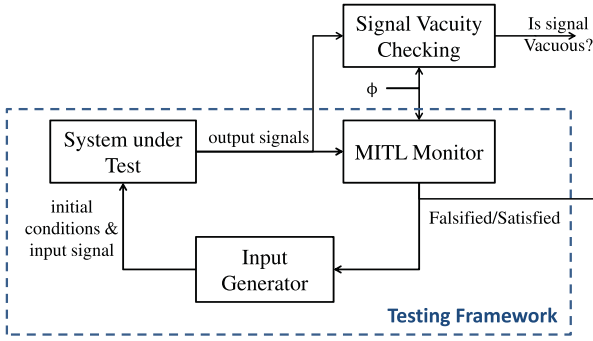


Fig. 6. Using signal vacuity checking to improve the confidence of an automatic test generation framework.




Fig. 7. The MITL SAT solver from Bersani et al. (2016) is used for debugging specifications.

Processor in the Loop (PiL), Hardware in the Loop (HiL), or real system. An example of a test generation technology that implements the architecture in Figure 6 is presented in Abbas et al. (2013). The system under test is simulated to generate an output trace. Then, a monitor checks the trace with respect to the specification and reports to the user whether the system trace satisfies or falsifies the specification (e.g., Maler and Nickovic (2004)). For each falsification, we will report to the user the falsifying trajectory to investigate the system for this error. Falsification-based approaches for CPS can help us find subtle bugs in industrial-size control systems (Jin et al. 2013). If after using stochastic-based testing and numerical analysis we cannot find any bugs, then we are more confident that the system works correctly. However, it will be concerning if the numerical analysis is mostly based on vacuous signals. This is because vacuous signals satisfy the specification for reasons other than what was originally intended. Signal vacuity checking is conducted in Figure 6 using Algorithm 6, and vacuous signals are reported to the user for further inspection. This will help users to focus their analysis on the part of the system that generates vacuous signals to prevent vacuous test generation.

6.2.1 Detecting Partially Covering Signals. A problem closely related to vacuity detection is the partial coverage problem. In this section, we show that Literal Occurrence Removal can be used for determining partially covering signals. Partially covering signals are the signals that not only satisfy the specification but also satisfy Literal Occurrence Removal mutation:

Definition 6.5 (Partially Covering Timed Trace (Signal)). Given an MITL (STL) formula ϕ in NNF, a timed trace μ (signal s) is partially covering if it satisfies the Literal Occurrence Removal mutation of ϕ .

This mutation is generated by repeatedly substituting the occurrences of literals with \perp , which is denoted by $\phi[l \leftarrow \perp]$ (see Definition 5.3). In Algorithm 7, we check whether the signal will satisfy the mutated specification ($\phi_i[l \leftarrow \perp]$). In the following, we prove that all satisfying signals are also partially covering signals:

THEOREM 6.6. *For all MITL formulas $\phi \in \Phi$ in NNF, if there exists a disjunction subformula in ϕ , then for all μ such that $\mu \models \phi$, it is always the case that there exists a literal $l \in \text{litOccur}(\phi)$ s.t. $\mu \models \phi[l \leftarrow \perp]$.*

The proof of Theorem 6.6 is provided in Appendix B. For any MITL (STL) specification ϕ that contains one or more disjunction operators (\vee) in NNF, any timed trace (signal) that satisfies ϕ will

also satisfy a mutation $\varphi[l \leftarrow \perp]$ for some literal occurrence l . Further, any specification that lacks a disjunction operator (\vee) in its NNF will not satisfy $\varphi[l \leftarrow \perp]$ for any literal occurrence l . That is, for formulas without any disjunction operator in NNF, we have $\varphi[l \leftarrow \perp] \equiv \perp$ since for any MITL formula φ , we have $\varphi \wedge \perp \equiv \perp$. If all satisfying signals are partially covering signals, what is the benefit of Algorithm 7? There are two applications for Algorithm 7 as follows.

First, Algorithm 7 can find which (how many) disjuncts are satisfied by the partially covering signal. This information can be used by the falsification technique to find the disjuncts/predicates that cause the formula satisfaction. Therefore, the falsification method will target the system behaviors corresponding to those predicates. As a result, Algorithm 7 can be used to improve the falsification method.

Second, Algorithm 7 can also be used for coverage analysis when falsification occurs. This is because with a slight modification, the dual of Algorithm 7 can help us to find the source of the falsification. According to Corollary B.2, for any φ in NNF, where φ has a conjunctive subformula of $\psi = \psi_1 \wedge \psi_2$, if $\mu \not\models \varphi$, then $\exists l \in \text{litOccur}(\psi)$ s.t. $\mu \not\models \varphi[l \leftarrow \top]$. Now, we can identify which conjunct of ψ contributes toward the falsification by substituting it by iteratively applying $\varphi[l \leftarrow \top]$. This can be better explained in the following example:

Example 6.7. Assume $\varphi = \diamond_{I_1}(a \wedge \diamond_{I_2}b)$ and a falsifying trace $\mu \not\models \varphi$ exists. Formula φ contains conjunction of $\psi = a \wedge \diamond_{I_2}b$ and $\text{litOccur}(\psi) = \{a, b\}$. We can substitute a and b with \top to find the main source of falsification of φ as follows:

- If $\mu \models \varphi[a \leftarrow \top]$, then $\mu \models \diamond_{I_1}(\diamond_{I_2}b)$, so a is the source of the problem.
- If $\mu \models \varphi[b \leftarrow \top]$, then $\mu \models \diamond_{I_1}(a)$, so b is the source of the problem.

As a result, the dual of Algorithm 7 can be used to debug a trace when the counter example is created using falsification methodologies (Abbas et al. 2013).

7 EXPERIMENTAL ANALYSIS

All three levels of the correctness analysis of MITL specifications need satisfiability checking as the underlying tool (Bersani et al. 2013). In validity checking, we simply check whether the specification and its negation are satisfiable. In general, in order to check whether $\varphi \models \psi$, we should check whether $\varphi \Rightarrow \psi$ is a tautology, that is, $\forall \mu, \mu \models \varphi \Rightarrow \psi$. This can be verified by checking whether $\neg(\varphi \Rightarrow \psi)$ is unsatisfiable. Recall that $\varphi \Rightarrow \psi$ is equivalent to $\neg\varphi \vee \psi$. So we have to check whether $\varphi \wedge \neg\psi$ is unsatisfiable to conclude that $\varphi \models \psi$. We use the above reasoning for redundancy checking as well as for vacuity checking. For redundancy checking of conjuncts at the root level, $\{\Phi \setminus \varphi_i\} \wedge \neg\varphi_i$ should be unsatisfiable, in order to conclude that $\{\Phi \setminus \varphi_i\} \models \varphi_i$. For vacuity checking, $\Phi \wedge \neg(\varphi_i[l \leftarrow \perp])$ should be unsatisfiable, in order to prove that $\Phi \models \varphi_i[l \leftarrow \perp]$.

7.1 MITL Satisfiability

The satisfiability problem of MITL is EXPSpace-complete (Alur et al. 1996). In order to check whether an MITL formula is satisfiable, we use two publicly available tools: qtlsolver⁷ and zot.⁸ The qtlsolver that we used translates MITL formulas into CLTL-over-clocks (Bersani et al. 2016, 2013). Constraint LTL (CLTL) is an extension of LTL where predicates are allowed to be assertions on the values of non-Boolean variables (Demri and D'Souza 2007). That is, in CLTL, we are allowed to define predicates using relational operators for variables over domains like \mathbb{N} and \mathbb{Z} . Although

⁷qtlsolver: A solver for checking satisfiability of Quantitative / Metric Interval Temporal Logic (MITL/QTL) over Reals. Available from <https://code.google.com/p/qtlsolver/>.

⁸The zot bounded model/satisfiability checker. Available from <https://code.google.com/p/zot/>.

Table 3. Incorrect Specifications from the Usability Study in Hoxha et al. (2015), Error Reported to the User by the Debugging Algorithm, and Algorithm Runtime (Formulas Have Been Translated from STL to MITL)

ϕ	Task #	MITL Specification Created by ViSPEC Users	Reporting the Errors	Sec.
ϕ_1	3	$\Diamond_{[0,30]}p_1 \wedge \Diamond_{[0,20]}p_1$	$\Diamond_{[0,30]}p_1$ is redundant	14
ϕ_2	3	$\Diamond_{[0,30]}(p_1 \Rightarrow \Box_{[0,20]}p_1)$	ϕ is a tautology	7
ϕ_3	10	$\Diamond_{[0,40]}(((p_1 \vee p_3) \Rightarrow \Diamond_{[0,20]}p_2) \wedge \Box_{[0,30]}p_1)$	ϕ is vacuous: $\phi \models \phi[p_3 \leftarrow \perp]$	39645
ϕ_4	4	$\Box_{[0,40]}p_1 \wedge \Box_{[0,40]}\Diamond_{[0,10]}p_1$	$\Box_{[0,40]}\Diamond_{[0,10]}p_1$ is redundant	29
ϕ_5	10	$\Diamond_{[0,40]}(p_1 \vee p_3) \wedge \Diamond_{[0,40]}p_2 \wedge \Diamond_{[0,40]}\Box_{[0,30]}p_1$	$\Diamond_{[0,40]}(p_1 \vee p_3)$ is redundant	126

satisfiability of CLTL in general is not decidable, some variants of it are decidable (Demri and D’Souza 2007).

CLTLoc (CLTL-over-clocks) is a variant of CLTL where the clock variables are the only arithmetic variables that are considered in the atomic constraints. It has been proven in Bersani et al. (2014) that CLTLoc is equivalent to timed automata (Clarke et al. 1999). Moreover, it can be polynomially reduced to decidable Satisfiability Modulo Theories that are solvable by many SMT solvers such as Z3.⁹ The satisfiability of CLTLoc is PSPACE-complete (Bersani et al. 2013) and the translation from MITL to CLTLoc in the worst case can be exponential (Bersani et al. 2016). Some restrictions must be imposed on the MITL formulas in order to use the qtsolver (Bersani et al. 2016). That is, the lower bound and upper bound for the intervals of MITL formulas should be integer values and the intervals are left/right closed. Therefore, we expect the values to be integer when we analyze MITL formulas. The high-level architecture of the MITL SAT solver, which we use to check the three issues, is provided in Figure 7.

7.2 Specification Debugging Results

We utilize the debugging algorithm on a set of specifications developed as part of a usability study for the evaluation of the ViSPEC tool (Hoxha et al. 2015). The usability study was conducted on two groups:

- (1) Group A: These are users who declared that they have little to no experience in working with requirements. The Group A cohort consists of 20 subjects from the academic community at Arizona State University. Most of the subjects have an engineering background.
- (2) Group B: These are users who declared that they have experience working with system requirements. Note that they do not necessarily have experience in writing requirements using formal logics. The Group B subject cohort was composed of 10 subjects from industry in the Phoenix metro area.

Each subject received a task list to complete. The list contained 10 tasks related to automotive system specifications. Each task asked the subject to formalize a natural language specification through ViSPEC and generate an STL specification. The task list is presented in Table 2. A detailed report on the accuracy of the users’ response to each natural language requirement is provided in Hoxha et al. (2015). Note that the specifications were preprocessed and transformed from the original STL formulas to MITL in order to run the debugging algorithm. For example, specification ϕ_3 in Table 3 originally in STL was $\phi_{3_{STL}} = \Diamond_{[0,40]}(((speed > 80) \Rightarrow \Diamond_{[0,20]}(rpm > 4000)) \wedge \Box_{[0,30]}(speed > 100))$. The STL predicate expressions $(speed > 80)$, $(rpm > 4000)$, $(speed > 100)$

⁹Microsoft Research, Z3: An efficient SMT solver. Available from <http://research.microsoft.com/en-us/um/redmond/projects/z3/>.

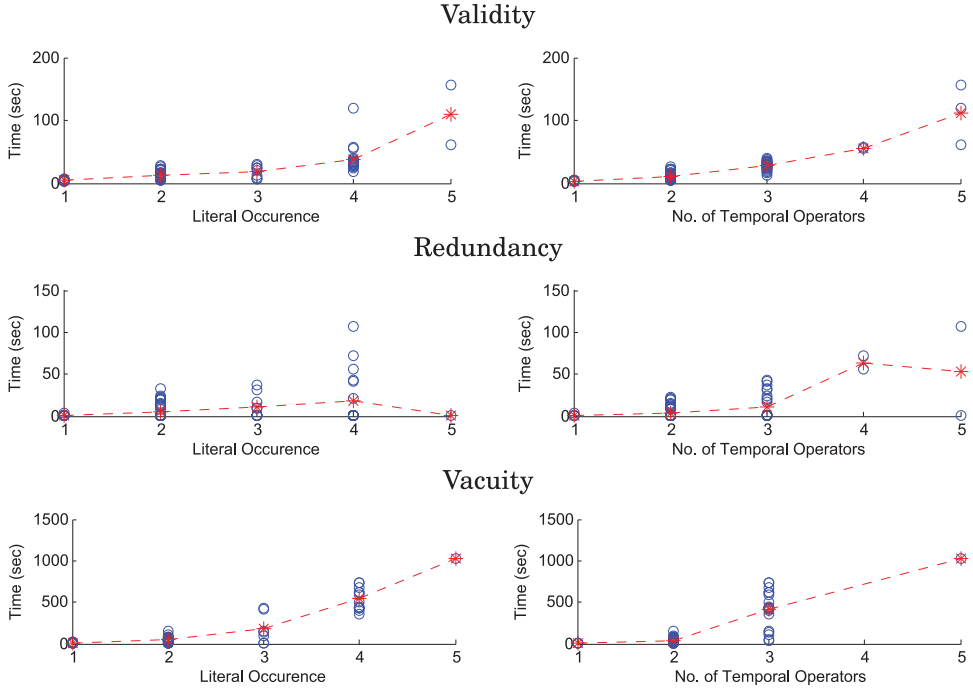


Fig. 8. Runtime overhead of the three stages of the debugging algorithm over user-submitted specifications. Timing results are presented over the number of literal occurrences and the number of temporal operators.

are mapped into atomic propositions with nonoverlapping predicates (Boolean functions) p_1, p_2, p_3 . The predicates p_1, p_2, p_3 correspond to the following STL representations: $p_1 \equiv \text{speed} > 100$, $p_2 \equiv \text{rpm} > 4,000$, and $p_3 \equiv 100 \geq \text{speed} > 80$. In Table 3, we present the common issues with the elicited specifications that our debugging algorithm detects. Note that validity, redundancy, and vacuity issues are present in the specifications listed. It should be noted that for specification ϕ_3 , although finding the error takes a significant amount of time, our algorithm can be used offline.

In Figure 8, we present the runtime overhead of the three-stage debugging algorithm over specifications collected in the usability study. In the first stage, 87 specifications go through validity checking. Five specifications fail the test and therefore they are immediately returned to the user. As a result, 82 specifications go through redundancy checking of conjunction in the root level,¹⁰ where nine fail the test. Lastly, 73 specifications go through vacuity checking, where five specifications have vacuity issues. The remaining 68 specifications passed the tests. Note that in the figure, two outlier data points are omitted from the vacuity subfigure for presentation purposes. The two cases were timed at 39,618sec and 17,421sec. In both cases, the runtime overhead was mainly because the zot software took hours to determine that the modified specification is unsatisfiable (both specifications were vacuous). The overall runtime of ϕ_3 in Table 3 is 39,645sec, which includes the runtime of validity and redundancy checking. The runtime overhead of vacuity checking of ϕ_3 can be reduced by half because, originally, in vacuity checking we run MITL satisfiability checking for all literal occurrences. In particular, ϕ_3 has four literal occurrences where for two cases zot took more than 19,500sec to determine that the modified specification is unsatisfiable. We can provide an option for early detection: stop and report as soon as an issue is found (the first unsatisfiability).

¹⁰In these experiments, we did not consider conjunctions in the lower-level subformulas for redundancy checking.

The circles in Figure 8 represent the timing performance in each test categorized by the number of literal occurrences and temporal operators. The asterisks represent the mean values and the dashed line is the linear interpolation between them. In general, we observe an increase in the average computation time as the number of literal occurrences and temporal operators increases. All the experimental results in Section 7 were performed on an Intel Xeon X5647 (2.993GHz) with 12GB RAM.

7.3 LTL Satisfiability

In the previous section, we mentioned that MITL satisfiability is a computationally hard problem. However, in practice, we know that LTL satisfiability is solvable faster than MITL satisfiability (Li et al. 2013). In this section, we consider how we can use the satisfiability of LTL formulas to decide about the satisfiability of MITL formulas. Consider the following fragments of MITL and LTL in NNF:

MITL(\Box): $\varphi ::= \top \mid \perp \mid p \mid \neg p \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \Box_I \varphi_1$

MITL(\Diamond): $\varphi ::= \top \mid \perp \mid p \mid \neg p \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \Diamond_I \varphi_1$

LTL(\Box): $\varphi ::= \top \mid \perp \mid p \mid \neg p \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \Box \varphi_1$

LTL(\Diamond): $\varphi ::= \top \mid \perp \mid p \mid \neg p \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid \Diamond \varphi_1$

In Appendix C, we prove that the satisfaction of a formula $\phi_M \in \text{MITL}(\Diamond)$ in NNF is related to the satisfaction of an LTL version of ϕ_M called $\phi_L \in \text{LTL}(\Diamond)$, where ϕ_L is identical to ϕ_M except that every interval I in ϕ_M is removed. For example, if $\phi_M = \Diamond_{[0,10]}(p \wedge q) \wedge \Diamond_{[0,10]}p$, then $\phi_L = \Diamond(p \wedge q) \wedge \Diamond p$. In essence, if ϕ_M is satisfiable, then ϕ_L is also satisfiable. Therefore, if ϕ_L is unsatisfiable, then ϕ_M is also unsatisfiable.

For the always (\Box) operator, satisfiability is the dual of the eventually operator (\Diamond). Assume that $\phi'_M \in \text{MITL}(\Box)$ contains only the \Box operator and $\phi'_L \in \text{LTL}(\Box)$ is the LTL version of ϕ'_M . If ϕ'_L is satisfiable, then ϕ'_M will also be satisfiable.

Based on the above discussion, if the specification that we intend to test/debug belongs to either category (fragment), MITL(\Diamond) or MITL(\Box), then we can check the satisfiability of its LTL version (ϕ_L) and decide according to the following:

THEOREM 7.1. *For any formula $\phi_M \in \text{MITL}(\Diamond)$ and $\phi'_M \in \text{MITL}(\Box)$ then:*

If $\phi_L \in \text{LTL}(\Diamond)$ is unsatisfiable, then ϕ_M is unsatisfiable.

If $\phi'_L \in \text{LTL}(\Box)$ is satisfiable, then ϕ'_M is satisfiable.

In these two cases, we do not need to run MITL SAT if otherwise we must apply MITL SAT, which means that we wasted effort by checking LTL SAT. However, since the runtime of LTL SAT is negligible, it will not drastically reduce the performance. As a result, LTL satisfiability checking is useful for validity testing. For redundancy checks, it may also be useful. For example, if we have a formula $\phi = \Diamond_{[0,10]}p \wedge \Box_{[0,20]}p$, we should check the satisfiability of $\phi' = \Box_{[0,10]}\neg p \wedge \Box_{[0,20]}p$ and $\phi'' = \Diamond_{[0,10]}p \wedge \Diamond_{[0,20]}\neg p$ for redundancy. Although the original formula ϕ does not belong to either MITL(\Diamond) or MITL(\Box), its modified NNF version will fit in these fragments and we may benefit by the usually faster LTL satisfiability for ϕ' and/or ϕ'' . For vacuity checking, we can use LTL satisfiability if, after manipulating/simplifying the original specification and creating the NNF version, we can categorize the resulting formula into the MITL(\Diamond) or the MITL(\Box) fragments (see Table 4).

We can check LTL satisfiability of the modified MITL specifications using existing methods and tools (Rozier and Vardi 2010). In our case, we used the NuSMV¹¹ tool with a similar encoding of LTL formulas as in Rozier and Vardi (2010). In Table 4, we compare the runtime overhead of MITL

¹¹NuSMV Version 2.6.0. Available from <http://nusmv.fbk.eu/>.

Table 4. Comparing the Runtime Overhead of MITL Satisfiability and LTL Satisfiability (in Seconds) for Some of the Specifications from ViSPEC's Usability Study

Test Phase	MITL Specification	MITL SAT	LTL SAT	MITL / LTL Runtime
Validity	$\Box_{[0,40]}(p_1 \Rightarrow \Box_{[0,10]}(p_1))$	4.154	0.047	88
Validity	$\Box_{[0,30]}(\neg p_1) \vee \Box_{[0,20]}(\neg p_1)$	3.418	0.0538	63
Validity	$\Box_{[0,40]}((\neg p_1 \wedge \neg p_3) \vee \Box_{[0,20]} \neg p_2 \vee \Box_{[0,30]} p_1)$	10.85	0.045	240
Validity	$\Box_{[0,40]}((p_1 \vee p_3) \Rightarrow \Box_{[0,20]}(p_2 \Rightarrow \Box_{[0,30]} p_1))$	15.406	0.0463	333
Vacuity	$\Box_{[0,40]}(p_1)$	1.71	0.0473	36
Vacuity	$\Box_{[0,40]}(p_1 \wedge \Box_{[0,10]}(p_1))$	3.727	0.044	84
Vacuity	$\Box_{[0,40]} p_1 \wedge \Box_{[0,30]}(p_4)$	5.77	0.0456	126
Vacuity	$\Box_{[0,40]} p_5 \wedge \Box_{[0,70]}(p_5)$	8.599	0.044	194

and LTL satisfiability checking. For the results of the usability study in Hoxha et al. (2015), we conduct validity and vacuity checking with the LTL satisfiability solver. We remark that in our results in Table 4, all the formulas belong to the MITL(\Box) fragment. Since we did not find any MITL(\Diamond) formula in our experiments where its LTL version is not satisfiable, we did not provide MITL(\Diamond) formulas in Table 4.

The first column of Table 4 provides the debugging test phase where we used the satisfiability checkers. The second column represents the MITL formulas that we tested using the SAT solver. We omit the LTL formulas from Table 4, since they are identical to MITL but do not contain timing intervals. The atomic propositions p_1, p_2, p_3, p_4, p_5 of the MITL formulas in Table 4 correspond to the following STL predicates: $p_1 \equiv \text{speed} > 100$, $p_2 \equiv \text{rpm} > 4,000$, $p_3 \equiv 100 \geq \text{speed} > 80$, $p_4 \equiv \text{rpm} > 3,000$, and $p_5 \equiv \text{speed} > 80$. The third and fourth columns represent the runtime overhead of satisfiability checking for MITL specifications and their corresponding LTL version. The last column represents the speedup of the LTL approach over the MITL approach. It can be seen that the LTL SAT solver (NuSMV) is about 30 to 300 times faster than the MITL SAT solver (zot). These results confirm that, when applicable, LTL SAT solvers outperform MITL SAT solvers in checking vacuity and validity issues in specifications. As a result, it is worth running LTL SAT before MITL SAT when it is possible.

7.4 Antecedent Failure Detection

To apply signal vacuity checking, we use the S-TaLiRo testing framework (Abbas et al. 2013; Hoxha et al. 2014b). S-TaLiRo is a MATLAB toolbox that uses stochastic optimization techniques to search for system inputs for Simulink models that falsify the safety requirements presented in MTL/STL (Abbas et al. 2013). The signal vacuity checking implemented in the S-TaLiRo tool is computationally efficient (PTIME). Its time complexity is proportional to the number of implication operations, the size of the formula, and the size of the signal (Fainekos et al. 2012).

In the following, we illustrate the vacuous signal detection process by using the Automatic Transmission (AT) model provided by Mathworks as a Simulink demo.¹² We introduced a few modifications to the model to make it compatible with the S-TaLiRo framework. Further details can be found in Hoxha et al. (2014a). S-TaLiRo calls the AT Simulink model in order to generate the output trajectories. The outputs contain two continuous-time real-valued signals: the speed of the engine ω (RPM) and the speed of the vehicle v . In addition, the outputs contain one continuous-time discrete-valued signal $gear$ with four possible values ($gear = 1, \dots, gear = 4$), which indicates

¹²Available at <http://www.mathworks.com/help/simulink/examples/modeling-an-automatic-transmission-controller.html>.

Table 5. Automatic Transmission Requirements Expressed in Natural Language and MITL from Hoxha et al. (2014a)

Req.	Natural Language	MITL Formula
ϕ_1^{AT}	There should be no transition from gear two to gear one and back to gear two in less than 2.5sec.	$\Box_{[0,27.5]}((g_2 \wedge \Diamond_{(0,0.04]}g_1) \Rightarrow \Box_{[0,2.5]}\neg g_2)$
ϕ_2^{AT}	After shifting into gear one, there should be no shift from gear one to any other gear within 2.5sec.	$\Box_{[0,27.5]}((\neg g_1 \wedge \Diamond_{(0,0.04]}g_1) \Rightarrow \Box_{[0,2.5]}g_1)$
ϕ_3^{AT}	If the ω is always less than 4,500, then the v cannot exceed 85 in less than 10sec.	$\Box_{[0,30]}(\omega \leq 4500) \Rightarrow \Box_{[0,10]}(v \leq 85)$
ϕ_4^{AT}	Within 10sec, v is less than 80 and from that point on, ω is always less than 4,500.	$\Diamond_{[0,10]}((v \leq 80) \Rightarrow \Box_{[0,30]}(\omega \leq 4500))$

Table 6. Reporting Signal Vacuity Issue for Each Mutated Formula

Requirement	Antecedent Failure Mutation	Vacuous Signals/All Signals
ϕ_1^{AT}	$\Box_{[0,27.5]}\neg(g_2 \wedge \Diamond_{(0,0.04]}g_1)$	1,989/2,000
ϕ_2^{AT}	$\Box_{[0,27.5]}\neg(\neg g_1 \wedge \Diamond_{(0,0.04]}g_1)$	1,994/2,000
ϕ_3^{AT}	$\neg\Box_{[0,30]}(\omega \leq 4500)$	60/214
ϕ_4^{AT}	$\Box_{[0,10]}\neg(v \leq 80)$	1,996/2,000

the current gear in the auto-transmission controller. S-TALiRo then monitors system trajectories with respect to the requirements provided in Table 5. There, in the MITL formulas, we use the shorthand g_i to indicate the gear value, i.e., $(gear = i) \equiv g_i$. The simulation time for the system is set to 30 seconds; therefore, we can use bounded MITL formulas for the requirements.

After testing the AT with S-TALiRo, we collected all the system trajectories. Then, we utilized the antecedent failure mutation on the specification to check signal vacuity (Algorithm 6) for each of the formulas that are provided in Table 5. We provide the antecedent failure specifications and the number of signals that satisfy them in Table 6. It can be seen in Table 6 that most of the system traces are vacuous signals where the antecedent is not satisfied. This helps the users to consider these issues and identify interesting test cases that can be used to initialize the system tester so that the antecedent is always satisfied. For an application of signal vacuity checking in falsification, see Dokhanchi et al. (2017).

8 CONCLUSION AND FUTURE WORK

We have presented a specification elicitation and debugging framework that can assist ViSPEC users to produce correct formal specifications. In particular, the debugging algorithm enables the detection of logical inconsistencies in MITL and STL specifications. Our algorithm improves the elicitation process by providing feedback to the users on validity, redundancy, and vacuity issues. In the future, the specification elicitation and debugging framework will be integrated in the ViSPEC tool to simplify MITL and STL specification development for verification of CPS. In addition, we considered vacuity detection with respect to signals. This enables improved analysis since some issues can only be detected when considering both the system and the specification. In the future, we will consider the feasibility of using vacuous signals to improve the counterexample generation process and system debugging using signal vacuity.

APPENDIX

A PROOF OF THEOREM 5.4

In order to show that Φ is inherently vacuous, we must show that if $\Phi \models \varphi_i[l \leftarrow \perp]$, then the mutated specification is equivalent to the original specification. In other words, we should show that

if $\Phi \models \varphi_i[l \leftarrow \perp]$, then $(\{\Phi \setminus \varphi_i\} \cup \varphi_i[l \leftarrow \perp]) \equiv \Phi$. If the mutated specification is equivalent to the original specification, then the original specification is vacuously satisfiable in any system. That is, the specification is *inherently vacuous* (Fisman et al. 2008; Chockler and Strichman 2009). We already know that if $\Phi \models \varphi_i[l \leftarrow \perp]$, then $\Phi \Rightarrow \varphi_i[l \leftarrow \perp]$ and trivially $\Phi \Rightarrow \varphi_i[l \leftarrow \perp] \cup \{\Phi \setminus \varphi_i\}$. Now we just need to prove the other direction. We need to prove that when φ_i is in NNF, then $\varphi_i[l \leftarrow \perp] \Rightarrow \varphi_i$. Since we replace only one specific literal occurrence of φ with \perp , the rest of the formula remains the same. Therefore, it should be noted that $\varphi_i[l \leftarrow \perp]$ does not modify any $l' \in \text{litOccur}(\varphi_i)$ where $l' \neq l$.

PROOF. We use structural induction to prove that $\varphi_i[l \leftarrow \perp] \Rightarrow \varphi_i$.

Base Case: $\varphi_i = l$ or $\varphi_i = l' \neq l$.

We know that $\perp \Rightarrow l$ and $l' \Rightarrow l'$. Therefore, $\varphi_i[l \leftarrow \perp] \Rightarrow \varphi_i$.

Induction Hypothesis: For any MITL φ_j in NNF, we have $\varphi_j[l \leftarrow \perp] \Rightarrow \varphi_j$ (or $\forall \varphi_j, \varphi_j[l \leftarrow \perp] \Rightarrow \varphi_j$).

Induction Step: We will separate the case into unary and binary operators.

Before providing the cases, we should review the *positively monotonic* operators (Kupferman and Vardi 2003). According to MITL semantics, $f \in \{\Box_I, \Diamond_I\}$ and $g \in \{\wedge, \vee\}$ are positively monotonic; i.e., for every MITL formulas φ_1 and φ_2 in NNF with $\varphi_1 \Rightarrow \varphi_2$, we have $f(\varphi_1) \Rightarrow f(\varphi_2)$. Also, for all MITL formulas φ' in NNF, we have $g(\varphi_1, \varphi') \Rightarrow g(\varphi_2, \varphi')$ and $g(\varphi', \varphi_1) \Rightarrow g(\varphi', \varphi_2)$.

Case 1: $\varphi_i = f(\varphi_j)$, where $f \in \{\Box_I, \Diamond_I\}$. Since f is positively monotonic, we have that $\varphi_j[l \leftarrow \perp] \Rightarrow \varphi_j$ implies $f(\varphi_j[l \leftarrow \perp]) \Rightarrow f(\varphi_j)$. Thus,

$$f(\varphi_j)[l \leftarrow \perp] = f(\varphi_j[l \leftarrow \perp]) \Rightarrow f(\varphi_j) = \varphi_i. \text{ As a result, } \varphi_i[l \leftarrow \perp] \Rightarrow \varphi_i.$$

Case 2: $\varphi_i = g(\varphi_{j_1}, \varphi_{j_2})$, where $g \in \{\wedge, \vee\}$. Since g is positively monotonic, we have that $\varphi_{j_1}[l \leftarrow \perp] \Rightarrow \varphi_{j_1}$, and $\varphi_{j_2}[l \leftarrow \perp] \Rightarrow \varphi_{j_2}$ implies

$$g(\varphi_{j_1}[l \leftarrow \perp], \varphi_{j_2}[l \leftarrow \perp]) \Rightarrow g(\varphi_{j_1}, \varphi_{j_2}). \text{ Thus, } g(\varphi_{j_1}, \varphi_{j_2})[l \leftarrow \perp] = g(\varphi_{j_1}[l \leftarrow \perp], \varphi_{j_2}[l \leftarrow \perp]) \Rightarrow g(\varphi_{j_1}, \varphi_{j_2}) = \varphi_i. \text{ As a result, } \varphi_i[l \leftarrow \perp] \Rightarrow \varphi_i.$$

Since $\varphi_i[l \leftarrow \perp] \Rightarrow \varphi_i$, we can have:

$$\{\Phi \setminus \varphi_i\} \cup \varphi_i[l \leftarrow \perp] \Rightarrow \{\Phi \setminus \varphi_i\} \cup \varphi_i, \text{ which is equivalent to}$$

$$\{\Phi \setminus \varphi_i\} \cup \varphi_i[l \leftarrow \perp] \Rightarrow \Phi. \quad \square$$

B PROOF OF THEOREM 6.6

In this section, we will prove that any MITL (STL) $\varphi \in \Phi$ that contains a disjunction operation (\vee) in NNF can be satisfied by partially covering signals. In other words, we will prove that any timed trace (signal) that satisfies φ will be considered as a partially covering timed trace (signal) according to Algorithm 7. Without loss of generality, we assume that both operands of disjunction are not constant. This is because if one of the operands is equivalent to \top or \perp , then the disjunction can be semantically removed as follows $\psi \vee \top \equiv \top$ or $\psi \vee \perp \equiv \psi$ for any MITL (STL) ψ .

Let us consider the partially covering timed trace (signal) returned by Algorithm 7. If there exist $\varphi_i \in \Phi$ and $l \in \text{litOccur}(\varphi_i)$ such that the timed trace μ satisfies $\varphi_i[l \leftarrow \perp]$, then μ will be reported as a partially covering timed trace (signal). Recall that we assume that Φ is a conjunction of MITL specifications according to Equation (1). We also assume that the conjunct $\varphi_i \in \Phi$ is the MITL subformula that contains the disjunction operation, namely, that $\psi = \psi_1 \vee \psi_2$ is a subformula of φ_i .

THEOREM B.1. Any timed trace μ that satisfies φ_i will satisfy $\varphi_i[l \leftarrow \perp]$ for some $l \in \text{litOccur}(\varphi_i)$.

PROOF. We have two cases for $\mu \models \varphi_i$ and $\psi \in \varphi_i$ where $\psi = \psi_1 \vee \psi_2$:

- (1) $\forall t, (\mu, t) \not\models \psi$: In this case, ψ does not affect the satisfaction of $\mu \models \varphi_i$. If we choose $l' \in \text{litOccur}(\psi)$, then $\psi[l' \leftarrow \perp]$ also does not affect the satisfaction of φ_i .

- (2) $\exists t$, s.t. $(\mu, t) \models \psi$: In this case, ψ affects the satisfaction of φ_i . So either $(\mu, t) \models \psi_1$ or $(\mu, t) \models \psi_2$. If $(\mu, t) \models \psi_1$, then we can choose $l' \in \text{litOccur}(\psi_2)$ and we have $(\mu, t) \models \psi_1 \vee \psi_2[l' \leftarrow \perp]$. Similarly, if $(\mu, t) \models \psi_2$, then we can choose $l'' \in \text{litOccur}(\psi_1)$ and we have $(\mu, t) \models \psi_1[l'' \leftarrow \perp] \vee \psi_2$. As a result, there exists some $l \in \text{litOccur}(\psi)$ where $(\mu, t) \models \psi[l \leftarrow \perp]$ and accordingly $\mu \models \varphi_i[l \leftarrow \perp]$.

Finally, $\forall \mu, \mu \models \varphi_i \exists l \in \text{litOccur}(\varphi_i)$ s.t. $\mu \models \varphi_i[l \leftarrow \perp]$, which means that μ is a partially covering timed trace (signal). \square

COROLLARY B.2. Assume that the conjunct $\varphi_j \in \Phi$ is the subformula that contains the conjunction operation in NNF, namely, that $\psi = \psi_1 \wedge \psi_2$ is a subformula of φ_j . Any timed trace μ that falsifies φ_j will falsify $\varphi_j[l \leftarrow \top]$ for some $l \in \text{litOccur}(\varphi_j)$.

C PROOFS OF THEOREM 7.1

We consider two MITL(\diamond, \square) fragments, denoted MITL(\square) and MITL(\diamond). In this proof, we assume that all formulas are in NNF. We also consider LTL(\diamond, \square) as the set of LTL formulas (with continuous semantics) that contains only \diamond and \square as temporal operators. In the following, we provide the continuous semantics of LTL(\diamond, \square) over traces with bounded duration. Semantics of LTL(\diamond, \square) over bounded timed traces can be defined as follows:

Definition C.1 (LTL(\diamond, \square) continuous semantics). Given a timed trace $\mu : [0, T] \rightarrow 2^{AP}$ and $t, t' \in \mathbb{R}$, and an LTL(\diamond, \square) formula ϕ , the satisfaction relation $(\mu, t) \models \phi$ for temporal operators is inductively defined:

- $(\mu, t) \models \diamond \phi_1$ if and only if $\exists t' \in [t, T]$ s.t. $(\mu, t') \models \phi_1$.
 $(\mu, t) \models \square \phi_1$ if and only if $\forall t' \in [t, T], (\mu, t') \models \phi_1$.

We will consider two LTL(\diamond, \square) fragments denoted LTL(\square) and LTL(\diamond). The syntax of MITL and LTL fragments are as presented in Section 7.3. We define the operator $[\phi]_{LTL}$, which can be applied to any MITL(\diamond, \square) formula and removes its interval constraints to create a new formula in LTL(\diamond, \square). For example, if $\phi = \diamond_{[0,10]}(p \wedge q) \wedge \diamond_{[0,10]}p \wedge \square_{[0,10]}q$, then $[\phi]_{LTL} = \diamond(p \wedge q) \wedge \diamond p \wedge \square q$. As a result, for any $\phi \in \text{MITL}(\diamond, \square)$ there exists a $\psi \in \text{LTL}(\diamond, \square)$ where $\psi = [\phi]_{LTL}$. For each MITL(\diamond, \square) formula ϕ , the language of ϕ denoted $L(\phi)$ is the set of all timed traces that satisfy ϕ : $\mu \models \phi$ if and only if $\mu \in L(\phi)$. Similarly, for any $\psi \in \text{LTL}(\diamond, \square)$, the language of ψ denoted $L(\psi)$ is the set of all timed traces that satisfy ψ : $\mu' \models \psi$ if and only if $\mu' \in L(\psi)$. Based on set theory, it is trivial to prove that $A \subseteq B$ and $C \subseteq D$ implies $A \cup C \subseteq B \cup D$ and $A \cap C \subseteq B \cap D$.

THEOREM C.2. For any formula $\varphi \in \text{MITL}(\diamond)$, and $t \in [0, T]$, we have $L_t(\varphi) \subseteq L_t([\varphi]_{LTL})$, where $L_t(\varphi) = \{\mu \mid (\mu, t) \models \varphi\}$. In other words, for every timed trace μ , we have that $(\mu, t) \models \varphi$ implies $(\mu, t) \models [\varphi]_{LTL}$.

PROOF. We use structural induction to prove that $L_t(\varphi) \subseteq L_t([\varphi]_{LTL})$.

Base Case: if $\varphi = \top, \perp, p, \neg p$, then $[\varphi]_{LTL} = \varphi$ and $L_t(\varphi) \subseteq L_t([\varphi]_{LTL})$.

Induction Hypothesis: We assume that there exist $\varphi_1, \varphi_2 \in \text{MITL}(\diamond)$ where for all $t \in [0, T]$, $L_t(\varphi_1) \subseteq L_t([\varphi_1]_{LTL})$ and $L_t(\varphi_2) \subseteq L_t([\varphi_2]_{LTL})$.

Case 1: For binary operators \wedge, \vee , we can use the union and intersection properties. In essence, for all formulas φ_1, φ_2 , we have $L_t(\varphi_1 \vee \varphi_2) = L_t(\varphi_1) \cup L_t(\varphi_2)$ and $L_t(\varphi_1 \wedge \varphi_2) = L_t(\varphi_1) \cap L_t(\varphi_2)$. According to the IH, $L_t(\varphi_1) \subseteq L_t([\varphi_1]_{LTL})$ and $L_t(\varphi_2) \subseteq L_t([\varphi_2]_{LTL})$; therefore, $L_t(\varphi_1) \cap L_t(\varphi_2) \subseteq L_t([\varphi_1]_{LTL}) \cap L_t([\varphi_2]_{LTL})$ and $L_t(\varphi_1) \cup L_t(\varphi_2) \subseteq L_t([\varphi_1]_{LTL}) \cup L_t([\varphi_2]_{LTL})$. As a result, $L_t(\varphi_1 \wedge \varphi_2) \subseteq L_t([\varphi_1]_{LTL} \wedge [\varphi_2]_{LTL}) = L_t([\varphi_1 \wedge \varphi_2]_{LTL})$, and $L_t(\varphi_1 \vee \varphi_2) \subseteq L_t([\varphi_1]_{LTL} \vee [\varphi_2]_{LTL}) = L_t([\varphi_1 \vee \varphi_2]_{LTL})$.

Case 2: For the temporal operator \diamond , we need to compare the semantics of $\text{MITL}(\diamond)$ and $\text{LTL}(\diamond)$.

Recall that

$(\mu, t) \models \diamond_I \varphi_1$ iff $\exists t' \in (t + I) \cap [0, T]$ s.t. $(\mu, t') \models \varphi_1$.

$(\mu, t) \models \diamond \varphi_1$ iff $\exists t' \in [t, T]$ s.t. $(\mu, t') \models \varphi_1$.

Recall that $t'' \in (t + I) \cap [0, T]$ implies $t'' \in [t, T]$ since the left bound of I is nonnegative.

According to the semantics, $\forall \mu. (\mu, t) \models \diamond_I \varphi_1$ implies

$\exists t' \in (t + I) \cap [0, T]$ s.t. $(\mu, t') \models \varphi_1$ implies

$\exists t' \in (t + I) \cap [0, T]$ s.t. $(\mu, t') \models [\varphi_1]_{\text{LTL}}$ according to IH ($L_{t'}(\varphi_1) \subseteq L_{t'}([\varphi_1]_{\text{LTL}})$).

If $\exists t' \in (t + I) \cap [0, T]$ s.t. $(\mu, t') \models [\varphi_1]_{\text{LTL}}$, then

$\exists t' \in [t, T]$ s.t. $(\mu, t') \models [\varphi_1]_{\text{LTL}}$ since $t' \in (t + I) \cap [0, T]$ implies $t' \in [t, T]$.

Moreover, $(\mu, t') \models [\varphi_1]_{\text{LTL}}$ implies that $(\mu, t) \models \diamond[\varphi_1]_{\text{LTL}} \equiv [\diamond \varphi_1]_{\text{LTL}}$.

As a result, $\forall \mu. (\mu, t) \models \diamond_I \varphi_1 \Rightarrow (\mu, t) \models [\diamond \varphi_1]_{\text{LTL}}$ so $L_t(\diamond_I \varphi_1) \subseteq L_t([\diamond \varphi_1]_{\text{LTL}})$. \square

If $\varphi \in \text{MITL}(\diamond)$, then $\overline{L_t([\varphi]_{\text{LTL}})} \subseteq \overline{L_t(\varphi)}$ (immediate from set theory). Thus, for all timed traces μ , $\mu \not\models [\varphi]_{\text{LTL}}$ implies that $\mu \not\models \varphi$.

COROLLARY C.3. For any $\varphi \in \text{MITL}(\diamond)$, if $[\varphi]_{\text{LTL}} \in \text{LTL}(\diamond)$ is unsatisfiable, then φ is unsatisfiable.

THEOREM C.4. For any formula $\varphi \in \text{MITL}(\square)$, and $t \in [0, T]$, we have $L_t([\varphi]_{\text{LTL}}) \subseteq L_t(\varphi)$, where $L_t(\varphi) = \{\mu \mid (\mu, t) \models \varphi\}$. In other words, $\forall \mu. (\mu, t) \models [\varphi]_{\text{LTL}} \Rightarrow (\mu, t) \models \varphi$.

PROOF. Similar to Theorem C.2, we can apply structural induction for the proof of Theorem C.4. \square

REFERENCES

- Houssam Abbas, Georgios Fainekos, Sriram Sankaranarayanan, Franjo Ivančić, and Aarti Gupta. 2013. Probabilistic temporal logic falsification of cyber-physical systems. *ACM Transactions on Embedded Computing Systems* 12, 2s, Article 95 (May 2013), 30 pages.
- Rajeev Alur, Costas Courcoubetis, Nicolas Halbwachs, Thomas A. Henzinger, Pei-Hsin Ho, Xavier Nicollin, Alfredo Olivero, Joseph Sifakis, and Sergio Yovine. 1995. The algorithmic analysis of hybrid systems. *Theoretical Computer Science* 138, 1 (1995), 3–34.
- Rajeev Alur, Tomás Feder, and Thomas A. Henzinger. 1996. The benefits of relaxing punctuality. *Journal of the ACM* 43, 1 (1996), 116–146.
- Glenn Ammons, David Mandelin, Rastislav Bodík, and James R. Larus. 2003. Debugging temporal specifications with concept analysis. In *Proceedings of the ACM SIGPLAN 2003 Conference on Programming Language Design and Implementation*. 182–195.
- Yashwanth Singh Rahul Annapureddy, Che Liu, Georgios E. Fainekos, and Sriram Sankaranarayanan. 2011. S-TaLiRo: A tool for temporal logic falsification for hybrid systems. In *Tools and Algorithms for the Construction and Analysis of Systems (LNCs)*, Vol. 6605. Springer, 254–257.
- Marco Autili, Paola Inverardi, and Patrizio Pelliccione. 2007. Graphical scenarios for specifying temporal properties: An automated approach. *Automated Software Engineering* 14, 3 (2007), 293–340.
- Thomas Ball and Orna Kupferman. 2008. Vacuity in testing. In *Proceedings of Tests and Proofs, Second International Conference (TAP'08)*. 4–17.
- Derek L. Beatty and Randal E. Bryant. 1994. Formally verifying a microprocessor using a simulation methodology. In *Proceedings of the 31th Design Automation Conference (DAC'94)*. 596–602.
- Ilan Beer, Shoham Ben-David, Cindy Eisner, and Yoav Rodeh. 2001. Efficient detection of vacuity in temporal model checking. *Formal Methods in System Design* 18, 2 (2001), 141–163.
- Shoham Ben-David, Fady Copt, Dana Fisman, and Sitvanit Ruah. 2015. Vacuity in practice: Temporal antecedent failure. *Formal Methods in System Design* 46, 1 (2015), 81–104.
- Marcello M. Bersani, Matteo Rossi, and Pierluigi San Pietro. 2014. A logical characterization of timed (non-)regular languages. In *Proceedings of the 39th International Symposium on Mathematical Foundations of Computer Science 2014, Part I (MFCS'14)*. 75–86.
- Marcello M. Bersani, Matteo Rossi, and Pierluigi San Pietro. 2016. A tool for deciding the satisfiability of continuous-time metric temporal logic. *Acta Informatica* 53, 2 (2016), 171–206.

- Marcello Maria Bersani, Matteo Rossi, and Pierluigi San Pietro. 2013. Deciding the satisfiability of MITL specifications. In *Proceedings of the 4th International Symposium on Games, Automata, Logics and Formal Verification (EPTCS'13)*, Vol. 119. Open Publishing Association, 64–78.
- Hana Chockler and Ofer Strichman. 2009. Before and after vacuity. *Formal Methods in System Design* 34, 1 (Feb. 2009), 37–58.
- Alessandro Cimatti, Marco Roveri, Viktor Schuppan, and Andrei Tchaltsev. 2008. Diagnostic information for realizability. In *Proceedings of the 9th International Conference on Verification, Model Checking, and Abstract Interpretation (VMCAI'08)*. 52–67.
- Edmund M. Clarke, Orna Grumberg, and Doron A. Peled. 1999. *Model Checking*. MIT Press, Cambridge, MA.
- Stéphane Demri and Deepak D'Souza. 2007. An automata-theoretic approach to constraint LTL. *Information and Computation* 205, 3 (2007), 380–415.
- Adel Dokhanchi, Bardh Hoxha, and Georgios E. Fainekos. 2015. Metric interval temporal logic specification elicitation and debugging. In *ACM/IEEE International Conference on Formal Methods and Models for Codesign (MEMOCODE'15)*. 70–79.
- Adel Dokhanchi, Shakiba Yaghoubi, Bardh Hoxha, and Georgios Fainekos. 2017. Vacuity aware falsification for MTL request-response specifications. In *Proceedings of the 13th IEEE Conference on Automation Science and Engineering (CASE'17)*.
- Alexandre Donze. 2010. Breach, A toolbox for verification and parameter synthesis of hybrid systems. In *Computer Aided Verification*. LNCS, Vol. 6174. Springer, 167–170.
- Matthew B. Dwyer, George S. Avrunin, and James C. Corbett. 1998. Property specification patterns for finite-state verification. In *Proceedings of the 2nd Workshop on Formal Methods in Software Practice (FMSP'98)*. ACM, 7–15.
- Rüdiger Ehlers and Vasumathi Raman. 2014. Low-effort specification debugging and analysis. In *Proceedings 3rd Workshop on Synthesis (SYNT'14)*. 117–133.
- Georgios Fainekos, Sriram Sankaranarayanan, Koichi Ueda, and Hakan Yazarel. 2012. Verification of automotive control applications using S-TaLiRo. In *Proceedings of the American Control Conference*.
- Georgios E. Fainekos. 2011. Revising temporal logic specifications for motion planning. In *IEEE International Conference on Robotics and Automation (ICRA'11)*. 40–45.
- Dana Fisman, Orna Kupferman, Sarai Sheinvald-Faragy, and Moshe Y. Vardi. 2008. A framework for inherent vacuity. In *Proceedings of the 4th International Haifa Verification Conference Hardware and Software: Verification and Testing (HVC'08)*. 7–22.
- Klaus Havelund and Grigore Rosu. 2004. Efficient monitoring of safety properties. *Journal on Software Tools for Technology Transfer* 6, 2 (2004), 158–173.
- Florian Horn, Wolfgang Thomas, Nico Wallmeier, and Martin Zimmermann. 2015. Optimal strategy synthesis for request-response games. *RAIRO - Theoretical Informatics and Applications* 49, 3 (2015), 179–203.
- Bardh Hoxha, Houssam Abbas, and Georgios Fainekos. 2014a. Benchmarks for temporal logic requirements for automotive systems. In *Proceedings of Applied Verification for Continuous and Hybrid Systems*.
- Bardh Hoxha, Hoang Bach, Houssam Abbas, Adel Dokhanchi, Yoshihiro Kobayashi, and Georgios Fainekos. 2014b. Towards formal specification visualization for testing and monitoring of cyber-physical systems. In *Proceedings of the International Workshop on Design and Implementation of Formal Tools and Systems*.
- Bardh Hoxha, Nikolaos Mavridis, and Georgios Fainekos. 2015. ViSPEC: A graphical tool for easy elicitation of MTL requirements. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*.
- Xiaoqing Jin, Alexandre Donze, Jyotirmoy Deshmukh, and Sanjit Seshia. 2013. Mining requirements from closed-loop control models. In *Hybrid Systems: Computation and Control*. ACM Press.
- James Kapinski, Jyotirmoy V. Deshmukh, Xiaoqing Jin, Hisahiro Ito, and Kenneth R. Butts. 2015. Simulation-guided approaches for verification of automotive powertrain control systems. In *Proceedings of the American Control Conference (ACC'15)*. 4086–4095.
- Kangjin Kim, Georgios E. Fainekos, and Sriram Sankaranarayanan. 2012. On the revision problem of specification automata. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'12)*. 5171–5176.
- Robert Könighofer, Georg Hofferek, and Roderick Bloem. 2013. Debugging formal specifications: A practical approach using model-based diagnosis and counterstrategies. *Journal on Software Tools for Technology Transfer* 15, 5–6 (2013), 563–583.
- Sascha Konrad and Betty H. C. Cheng. 2005. Real-time specification patterns. In *Proceedings of the 27th International Conference on Software Engineering (ICSE'05)*. ACM, 372–381.
- Ron Koymans. 1990. Specifying real-time properties with metric temporal logic. *Real-Time Systems* 2, 4 (1990), 255–299.
- Orna Kupferman and Moshe Y. Vardi. 2003. Vacuity detection in temporal model checking. *Journal on Software Tools for Technology Transfer* 4, 2 (2003), 224–233.
- Jianwen Li, Lijun Zhang, Geguang Pu, Moshe Y. Vardi, and Jifeng He. 2013. LTL satisfiability checking revisited. In *2013 20th International Symposium on Temporal Representation and Reasoning*. 91–98.

- Oded Maler and Dejan Nickovic. 2004. Monitoring temporal properties of continuous signals. In *Proceedings of FORMATS-FTRTFT (LNCS)*, Vol. 3253. 152–166.
- Roland Meyer, Johannes Faber, Jochen Hoenicke, and Andrey Rybalchenko. 2008. Model checking duration calculus: A practical approach. *Formal Aspects of Computing* 20, 4 (2008), 481–505.
- Amalinda Post, Jochen Hoenicke, and Andreas Podelski. 2011. Vacuous real-time requirements. In *19th IEEE International Requirements Engineering Conference (RE’11)*. 153–162.
- Vasumathi Raman and Hadas Kress-Gazit. 2011. Analyzing unsynthesizable specifications for high-level robot behavior using LTLMoP. In *Proceedings of the 23rd International Conference on Computer Aided Verification (CAV’11)*. 663–668.
- Kristin Y. Rozier and Moshe Y. Vardi. 2010. LTL satisfiability checking. *Journal on Software Tools for Technology Transfer* 12, 2 (2010), 123–137.
- Hengyi Yang, Bardh Hoxha, and Georgios Fainekos. 2012. Querying parametric temporal logic properties on embedded systems. In *Testing Software and Systems*. Springer, 136–151.
- Pengcheng Zhang, Bixin Li, and Lars Grunske. 2010. Timed property sequence chart. *Journal of Systems and Software* 83, 3 (2010), 371–390.

Received January 2016; revised April 2017; accepted September 2017