Property specification patterns at work: verification and inconsistency explanation

Massimo Narizzano[1] · Luca Pulina[2] · Armando Tacchella[1] · Simone Vuotto[1,2]

Abstract
Property specification patterns (PSPs) have been proposed to ease the formalization of requirements, yet enable automated verification thereof. In particular, the internal consistency of specifications written with PSPs can be checked automatically with the use of, for example, linear temporal logic (LTL) satisfiability solvers. However, for most practical applications, the expressiveness of PSPs is too restricted to enable writing useful requirement specifications, and proving that a set of requirements is inconsistent can be worthless unless a minimal set of conflicting requirements is extracted to help designers to correct a wrong specification. In this paper, we extend PSPs by considering Boolean as well as atomic numerical assertions, we contribute an encoding from extended PSPs to LTL formulas, and we present an algorithm computing inconsistency explanations, i.e., irreducible inconsistent subsets of the original set of requirements. Our extension enables us to reason about the internal consistency of functional requirements which would not be captured by basic PSPs. Experimental results demonstrate that our approach can check and explain (in)consistencies in specifications with nearly two thousand requirements generated using a probabilistic model, and that it enables effective handling of real-world case studies.

Keywords Consistency of requirements · Property specifications patterns · LTL satisfiability checking · Inconsistency explanation

1 Introduction

In the context of safety- and security-critical cyber-physical systems (CPSs), checking the sanity of functional requirements is an important, yet challenging task. Requirements written in natural language call for time-consuming and error-prone manual reviews, whereas enabling automated sanity verification often requires overburdening formalizations. Given the increasing pervasiveness of CPSs, their strin-

gent time-to-market and product budget constraints, practical solutions to enable automated verification of requirements are in order. Property specification patterns (PSPs) [18] offer a viable path toward this target. PSPs are a collection of parameterizable, high-level, formalism-independent specification abstractions, originally developed to capture recurring solutions to the needs of requirement engineering. Each pattern can be directly encoded in a formal specification language, such as linear temporal logic (LTL) [39], computational tree logic (CTL) [10], or graphical interval logic (GIL) [14]. Because of their features, PSPs may ease the burden of formalizing requirements, yet enable verification of their sanity using current state-of-the-art automated reasoning tools—see, e.g., [9,21,25,28,46].

Sanity checking of requirements may consist of three parts: redundancy (vacuity) checking, completeness checking and consistency checking [3]. A specification is satisfied vacuously in a model if it is satisfied in some non-interesting way; borrowing the example from [43], the LTL specification "every request is eventually followed by a grant" is satisfied vacuously in a model with no requests. Vacuity checking can also be performed without the need of a model, and in

✉ Simone Vuotto
svuotto@uniss.it

Massimo Narizzano
massimo.narizzano@unige.it

Luca Pulina
lpulina@uniss.it

Armando Tacchella
armando.tacchella@unige.it

[1] DIBRIS, University of Genoa, Viale Causa 13, 16145 Genoa, Italy

[2] Chemistry and Pharmacy Department, University of Sassari, Via Vienna 2, Sassari, Italy

this case it is known as *inherent* vacuity checking [19,44]. Completeness checking is equivalent to verify if the set of requirements covers all reasonable behaviors of a system. Completeness can be checked in combination with a system model, but in [3] a proposal for model-free completeness checking is also presented. Finally, requirements consistency is about checking whether a *real* system can be implemented from a given set of requirements. Therefore, two types of check [44] are possible: (*i*) *realizability*, i.e., testing whether there is an *open* system that satisfies all the properties in the set [40], and (*ii*) *satisfiability*, i.e., testing whether there is a *closed* system that satisfies all the properties in the set. Satisfiability checking ensures that the behavioral description of a system is internally consistent and neither over- or under-constrained. If a property is either valid, or unsatisfiable, this must be due to an error. Even if the satisfiability test is weaker than the realizability test, its importance is widely recognized [44]. In this paper, we restrict our attention to sanity checking as satisfiability checking. We speak of (internal) consistency of requirements written using PSPs having in mind that PSPs can be translated to LTL formulas whose satisfiability can be checked using methods and tools available in the literature—see, e.g., [23,33,46,47] for tableau-based methods and [26,27,42–44] for methods based on automata-theoretic approaches.

The original formulation of PSPs caters for temporal structure over Boolean variables, but for most practical applications such expressiveness is too restricted. This is the case of the embedded controller for robotic manipulators that is under development in the context of the EU project CERBERO [35][1] and provides the main motivation for this work. As an example, consider the following statement: "The angle of joint1 shall never be greater than 170 degrees." This requirement imposes a safety threshold related to some joint of the manipulator (*joint1*) with respect to physically realizable poses, yet it cannot be expressed as a PSP unless we add atomic numerical assertions in some constraint system $\mathcal{D}$. We call constraint PSP, or PSP($\mathcal{D}$) for short, a pattern which has the same structure of a PSP, but contains atomic propositions from $\mathcal{D}$. For instance, using PSP($\mathbb{R}, <, =$) we can rewrite the above requirement as a *universality* pattern: "Globally, it is always the case that $\theta_1 < 170$ holds," where θ_1 is the numerical signal (variable) for the angle of *joint1*. In principle, automated reasoning about constraint PSPs can be performed in constraint linear temporal logic, i.e., LTL extended with atomic assertions from a constraint system [12]: in our example above, the encoding would be simply $\Box(\theta_1 < 170)$. Unfortunately, this approach does not always lend itself to a practical solution, because constraint linear

temporal logic is undecidable in general [11]. Restrictions on $\mathcal{D}$ may restore decidability [12], but they introduce limitations in the expressiveness of the corresponding PSPs. In this paper, we propose a solution which ensures that automated verification of consistency is feasible, yet enables PSPs mixing both Boolean variables and (constrained) numerical signals. Our approach enables us to capture many specifications of practical interest and to pick a verification procedure from the relatively large pool of automated reasoning systems currently available for LTL. In particular, we restrict our attention to constraint systems of the form $(\mathbb{R}, <, =)$ and atomic propositions of the form $x < c$ or $x = c$, where $x \in \mathbb{R}$ is a variable and $c \in \mathbb{R}$ is a constant value. In the following, we write $\mathcal{D}_C$ to denote such restriction.

Knowing that a set of requirements written with PSPs($\mathcal{D}_C$) is (in)consistent is only the first step in writing a correct specification. In case of inconsistent requirements, obtaining a *minimal* set of such requirements would be desirable to help designers avoid manual checks to pinpoint problems in a specification. The problem of finding minimal unsatisfiable subsets, or *inconsistency explanations*, has been the subject of some attention, e.g., in propositional satisfiability and constraint programming. The algorithms to be found in the literature can be either domain specific—see, e.g., [5,30]—or domain independent—see, e.g., [22]. They can be further divided into algorithms that find only one inconsistent subset or all inconsistent subsets. To the best of our knowledge, most approaches use a general-purpose algorithm, also known as the "deletion algorithm." Recently, some special-purpose algorithms have been proposed: (*i*) PROCMINE [1] uses a tableau-based solver to obtain an initial subset from an unsatisfiable set of LTL formulas and then applies deletion-based minimization to that subset; (*ii*) PLTL- MUP [20] uses a method based on ordered binary decision diagrams to find inconsistent subsets; (*iii*) TRP++UC [45] uses resolution graphs to extract minimal unsatisfiable subsets of requirements by modifying the algorithm implemented in the temporal resolution-based solver TRP++.

Since for practical reasons in requirement engineering it is better to have a quick turnaround time rather than a complete answer, we present a method to look for inconsistencies in an incremental fashion, i.e., stopping the search once at least one (minimal) inconsistency subset is found. In particular, given a set of inconsistent requirements, we extract a minimal (irreducible) subset from them that it is still inconsistent. The set is guaranteed to be minimal in the sense that, if we remove one of the elements, the remaining set becomes consistent.

Overall, our contribution can be summarized as follows:

– We extend basic PSPs over the constraint system $\mathcal{D}_C$.
– We provide an encoding from any PSP($\mathcal{D}_C$) into a corresponding LTL formula.

---

[1] Cross-layer modEl-based fRamework for multi-oBjective dEsign of Reconfigurable systems in unceRtain hybRid envirOnments—http://www.cerbero-h2020.eu/.

– We present a tool[2] based on state-of-the-art decision procedures and model checkers to automatically analyze requirements expressed as PSPs($\mathcal{D}_C$).

– We propose algorithms devoted to extract minimal subsets of inconsistent requirements, and we implement them in the tool mentioned above.

– We implement a generator of artificial requirements expressed as PSPs($\mathcal{D}_C$); the generator takes a set of parameters in input and emits a collection of PSPs according to a parameterized probability model.

– Using our generator, we run an extensive experimental evaluation aimed at understanding (i) which automated reasoning tool is best at handling set of requirements as PSPs($\mathcal{D}_C$), and (ii) whether our approach is scalable.

– Finally, we analyze the specification of the embedded controller to be dealt with in the context of CERBERO project, experimenting also with the addition of faulty requirements.

Verification and inconsistency explanation of requirements written in PSP($\mathcal{D}_C$) are carried out using tools and techniques available in the literature [25,43,44]. With those, we demonstrate the scalability of our approach by checking the consistency of up to 1920 requirements, featuring 160 variables and up to 8 different constant values appearing in atomic assertions, within less than 500 CPU s. A total of 75 requirements about the embedded controller for the CERBERO project is checked in a matter of seconds, even without resorting to the best tool among those we consider. This paper is based on and extends the one presented at the NASA Formal Method Conference [37]. The additional material relates to (i) algorithms for inconsistency explanation, including their experimental evaluation, (ii) proofs of results that were only stated in [37], and (iii) the experimental analysis of the tableau-based tool LEVIATHAN[7].

The rest of the paper is organized as follows. Section 2 contains some basic concepts on LTL, PSPs and some related work. In Sect. 3, we present the extension of basic PSPs over $\mathcal{D}_C$ and the related encoding to LTL, while in Sect. 4 we present inconsistency explanation algorithms. In Sects. 5 and 6, we report the results of the experimental analysis concerning the scalability and the case study on the embedded controller, respectively. We conclude the paper in Sect. 7 with some final remarks.

2 Background and related work

LTL syntax and semantics Linear temporal logic (LTL) [38] formulas are built on a finite set $Prop$ of atomic propositions

as follows:

$$\phi = p \mid \neg\phi_1 \mid \phi_1 \vee \phi_2 \mid \mathcal{X}\,\phi_1 \mid \phi_1\,\mathcal{U}\,\phi_2 \mid (\phi),$$

where $p \in Prop, \phi, \phi_1, \phi_2$ are LTL formulas, $\mathcal{X}$ is the "next" operator and $\mathcal{U}$ is the "until" operator. In the following, unless specified otherwise using parentheses, unary operators have higher precedence than binary operators. An LTL formula is interpreted over a *computation*, i.e., a function $\pi : \mathbb{N} \to 2^{Prop}$ which assigns truth values to the elements of $Prop$ at each time instant (natural number). For a computation π and a time instant $i \in \mathbb{N}$:

– $\pi, i \models p$ for $p \in Prop$ iff $p \in \pi(i)$
– $\pi, i \models \neg\alpha$ iff $\pi, i \not\models \alpha$
– $\pi, i \models (\alpha \vee \beta)$ iff $\pi, i \models \alpha$ or $\pi, i \models \beta$
– $\pi, i \models \mathcal{X}\,\alpha$ iff $\pi, i + 1 \models \alpha$
– $\pi, i \models \alpha\,\mathcal{U}\,\beta$ iff for some $j \geq i$, we have $\pi, j \models \beta$ and for all $k, i \leq k < j$ we have $\pi, k \models \alpha$.

We say that π *satisfies* a formula ϕ, denoted $\pi \models \phi$, iff $\pi, 0 \models \phi$. If $\pi \models \phi$ for every π, then ϕ is *true* and we write $\models \phi$.

We consider other Boolean connectives like "$\wedge$" and "$\to$" with the usual meaning, and we abbreviate $p \vee \neg p$ as $\top$, $p \wedge \neg p$ as $\bot$. We introduce $\Diamond \phi$ ("eventually") to denote $\top\,\mathcal{U}\,\phi$ and $\Box \phi$ ("always") to denote $\neg\Diamond\neg\phi$. Finally, some of the PSPs use the "weak until" operator defined as $\alpha\,\mathcal{W}\,\beta = \Box\alpha \vee (\alpha\,\mathcal{U}\,\beta)$.

LTL satisfiability Among various approaches to decide LTL satisfiability, reduction to model checking was proposed in [42] to check the consistency of requirements expressed as LTL formulas. Given a formula ϕ over a set $Prop$ of atomic propositions, a *universal* model M can be constructed. Intuitively, a universal model encodes all the possible computations over $Prop$ as (infinite) traces, and therefore ϕ is satisfiable precisely when M does not satisfy $\neg\phi$. In [44], a first improvement over this basic strategy is presented together with the tool PANDA,[3] whereas in [27] an algorithm based on automata construction is proposed to enhance performances even further—the approach is implemented in a tool called AALTA. Further studies along this direction include [26] and [25]. In the latter, a portfolio LTL satisfiability solver called POLSAT is proposed to run different techniques in parallel and return the result of the first one to terminate successfully.

Property specification patterns (PSPs) The original proposal of PSPs is to be found in [18]. They are meant to describe the structure of systems' behaviors and provide expressions of such behaviors in a range of common formalisms. An example of a PSP is shown in Fig. 1—with some parts omitted for

> **Response**
>
> Describe cause-effect relationships between a pair of events/states. An occurrence of the first, the cause, must be followed by an occurrence of the second, the effect. Also known as Follows and Leads-to.
>
> **Structured English Grammar**
> *It is always the case that if P holds, then S eventually holds.*
>
> **LTL Mappings**
>
> | | |
> |---|---|
> | Globally | $\Box (P \to \Diamond S)$ |
> | Before R | $\Diamond R \to (P \to (\overline{R} \; \mathcal{U} \; (S \wedge \overline{R}))) \; \mathcal{U} \; R$ |
> | After Q | $\Box (Q \to \Box (P \to \Diamond S))$ |
> | Between Q and R | $\Box ((Q \wedge \overline{R} \wedge \Diamond R) \to (P \to (\overline{R} \; \mathcal{U} \; (S \wedge \overline{R}))) \; \mathcal{U} \; R)$ |
> | After Q until R | $\Box (Q \wedge \overline{R} \to ((P \to (\overline{R} \; \mathcal{U} \; (S \wedge \overline{R}))) \; \mathcal{W} \; R)$ |
>
> **Example**
> *It is always the case that if object_detected holds , then moving_to_target eventually holds.*

Fig. 1 Response pattern ($\overline{\alpha}$ stands for $\neg\alpha$)

the sake of readability.[4] A pattern is comprised of a *Name* (response in Fig. 1), an (informal) statement describing the behavior captured by the pattern and a (structured English) statement [24] that should be used to express requirements. The LTL mappings corresponding to different declinations of the pattern are also given, where capital letters (P, S, T, etc.) stand for Boolean states/events. In more detail, a PSP is composed of two parts: (*i*) the *scope* and (*ii*) the *body*. The *scope* is the extent of the program execution over which the pattern must hold, and there are five scopes allowed: *Globally*, to span the entire scope execution; *Before*, to span execution up to a state/event; *After*, to span execution after a state/event; *Between*, to cover the part of execution from one state/event to another one; and *After-until*, where the first part of the pattern continues even if the second state/event never happens. For state-delimited scopes, the interval in which the property is evaluated is closed at the left and open at the right end. The *body* of a pattern describes the behavior that we want to specify. In [18], bodies are categorized in *occurrence* and *order* patterns. Occurrence patterns require states/events to occur or not to occur. Examples of such bodies are *Absence*, where a given state/event must not occur within a scope, and its opposite *Existence*. Order patterns constrain the order of the states/events. Examples of such patterns are *Precedence*, where a state/event must always precede another state/event, and *Response*, where a state/event must always be followed by another state/event within the scope. Moreover, we included the *Invariant* pattern introduced in [41], and dictating that a state/event must occur whenever another state/event occurs. Combining scopes and bodies, we can construct 55 different types of patterns.

Inconsistency explanation Usually, inconsistency in a set of requirements is best explained in terms of minimal subsets of requirements exposing the core issues within the specification. The literature does not provide a consistent naming of such cores, and the terms *minimal inconsistency subset* (*MIS*) [6], *minimal unsatisfiable subset* [5] (*MUS*), *minimal unsatisfiable core* [30] (*MUC*) and also high-level MUC (HLMUC) [36] are introduced to refer to the same concept—in the following, and throughout the paper, we denote with MUC a minimal set of inconsistent requirements. Algorithms for finding MUCs can be divided in two basic groups: (*i*) those focusing on the extraction of a single MUC and (*ii*) those focusing on the extraction of all MUCs. These techniques can be further divided into domain specific, i.e., targeting specific domains such as propositional satisfiability [4], and general purpose, i.e., high-level algorithms that can be applied to any domain provided that a consistency checking procedure exists for that domain [17]. The most basic general-purpose solution for computing a single MUC out of a set of logical constraints consists of iteratively removing constraints from an initial set. At each step, the set of constraints represents an over-approximation of the MUC. This solution is referred to as the deletion-based approach [2,8,13,17]. Given a set R of n constraints, the deletion-based approach calls the consistency checker exactly n times. When examining the ith constraint, if $R \setminus \{r_i\}$ remains inconsistent, then there is a MUC that does not include r_i, and r_i can be removed; otherwise, r_i must be part of the MUC. This approach is guaranteed to produce a set $M \subseteq R$ such that, if a single requirement is eliminated from M, then M becomes consistent. However, the approach does not guarantee that another MUC $M' \subseteq R$ such that $|M'| \leq |M|$ may not exist. The majority of the algorithms presented in the literature are domain specific [5,29,30,34] and, to the best of our knowledge, no specific approach that works for LTL has been proposed so far. Extraction of all MUCs has received some attention, also because retrieving MUCs of minimal size can be done simply by enumerating all MUCs. Finding all the MUCs of a set of constraints R in a naive way amounts to check the consistency of all the elements of the power set 2^R, but this is clearly untenable in real-world applications. In [29], the power set of requirements is implicitly considered as follows. Given a set of requirements R, if $R' \subseteq R$ is inconsistent, every $R'' \supset R'$ and $R'' \subset R$ is also inconsistent. Furthermore, if $R' \subseteq R$ is consistent, every $R'' \subset R'$ is consistent too. This algorithm can be modified to find a single MUC by stopping it to the first MUC extracted.

Related work In [31], the framework *Property Specification Pattern Wizard (PSP-Wizard)* is presented. Its purpose is the machine-assisted definition of temporal formulas capturing pattern-based system properties. PSP-Wizard offers a translation into LTL of the patterns encoded in the tool, but it is meant to aid specification, rather than support consis-

---

[4] We omitted aspects which are not relevant for our work, e.g., translations to other logics like CTL [18]. The full list of PSPs considered in this paper and their mapping to LTL and other logics is available at http://patterns.projects.cis.ksu.edu/.

tency checking, and it cannot deal with numerical signals. In [24], an extension is presented to deal with real-time specifications, together with mappings to metric temporal logic (MTL), timed computational tree logic (TCTL) and real-time graphical interval logic (RTGIL). Even if this work is not directly connected with ours, it is worth mentioning it since their structured English grammar for patterns is at the basis of our formalism. The work in [24] also provided inspiration to a recent set of works [15,16] about a tool, called VI-Spec, to assist the analyst in the elicitation and debugging of formal specifications. VI-Spec lets the user specify requirements through a graphical user interface, translates them to MITL formulas and then supports debugging of the specification using run-time verification techniques. VI-Spec embodies an approach similar to ours to deal with numerical signals by translating inequalities to sets of Boolean variables. However, VI-Spec differs from our work in several aspects, most notably the fact that it performs debugging rather than consistency checking, so the behavior of each signal over time must be known. Also, VI-Spec handles only inequalities and does not deal with sets of requirements written using PSPs.

3 Constraint property specification patterns

Let us start by defining a *constraint system* $\mathcal{D}$ as a tuple $\mathcal{D} = (D, R_1, \ldots, R_n, \mathcal{I})$, where D is a non-empty set called *domain*, and each R_i is a predicate symbol of arity a_i, with $\mathcal{I}(R_i) \subseteq D^{a_i}$ being its interpretation. Given a finite set of variables X and a finite set of constants C such that $C \cap X = \emptyset$, a *term* is a member of the set $T = C \cup X$; an (atomic) $\mathcal{D}$-*constraint* over a set of terms is of the form $R_i(t_1, \ldots, t_{a_i})$ for some $1 \le i \le n$ and $t_j \in T$ for all $1 \le j \le a_i$ which we call *constraint* when $\mathcal{D}$ is understood from the context. We define *linear temporal logic modulo constraints*—LTL($\mathcal{D}$) for short—as an extension of LTL with additional atomic constraints. Given a set of Boolean propositions $Prop$, a constraint system $\mathcal{D} = (D, R_1, \ldots, R_n, \mathcal{I})$ and a set of terms $T = C \cup X$, an LTL($\mathcal{D}$) formula is defined as:

$$\phi = p \mid R_i(t_1, \ldots, t_{a_i}) \mid \neg\phi_1 \mid \phi_1 \vee \phi_2 \mid \mathcal{X}\,\phi_1 \mid \phi_1\,\mathcal{U}\,\phi_2 \mid (\phi),$$

where $p \in Prop$, ϕ, ϕ_1, ϕ_2 are LTL($\mathcal{D}$) formulas and $R_i(\cdot)$ with $1 \le i \le n$ is an atomic $\mathcal{D}$-constraint. Additional Boolean and temporal operators are defined as in LTL with the same intended meaning. Notice that the set of LTL($\mathcal{D}$) formulas is a (strict) subset of those in constraint linear temporal logic—CLTL($\mathcal{D}$) for short—as defined, e.g., in [12]. LTL($\mathcal{D}$) formulas are interpreted over computations of the form $\pi : \mathbb{N} \to 2^{Prop}$ plus additional *evaluations* of the form $\nu : T \times \mathbb{N} \to D$ such that, for all $i \in \mathbb{N}$, $\nu(c, i) = \nu(c) \in D$ for all $c \in C$, whereas $\nu(x, i) \in D$ for all $x \in X$. In other words, the function ν associates with constants $c \in C$ a value

$\nu(c)$ that does not change in time, and to variables $x \in X$ a value $\nu(x, i)$ that possibly changes at each time instant $i \in \mathbb{N}$. LTL semantics is extended to LTL($\mathcal{D}$) by handling constraints:

$$\pi, \nu, j \models_{\mathcal{D}} R_i(t_1, \ldots, t_{a_i}) \text{iff} (\nu(t_1, j), \ldots, \nu(t_{a_i}, j)) \in \mathcal{I}(R_i).$$

We say that π and ν *satisfy* a formula ϕ, denoted $\pi, \nu \models_{\mathcal{D}} \phi$, iff $\pi, \nu, 0 \models \phi$. A formula ϕ is *satisfiable* as long as there exist a computation π and a valuation ν such that $\pi, \nu \models_{\mathcal{D}} \phi$. We further restrict our attention to the constraint system $D_C = (\mathbb{R}, <, =, \mathcal{I})$, with atomic constraints of the form $x < c$ and $x = c$, where c is a constant corresponding to some real number—hereafter we abuse notation and write $c \in \mathbb{R}$ instead of $\nu(c) \in \mathbb{R}$—and the interpretation $\mathcal{I}$ of the predicates "$<$" and "$=$" is the usual one. While CLTL($\mathcal{D}$) is undecidable in general [11,12], LTL($\mathcal{D}_C$) is decidable since, as we show in this paper, it can be reduced to LTL satisfiability.

We introduce the concept of *constraint property specification pattern*, denoted PSP($\mathcal{D}$), to deal with specifications containing Boolean variables as well as atoms from a constraint system $\mathcal{D}$. In particular, a PSP($\mathcal{D}_C$) features only Boolean atoms and atomic constraints of the form $x < c$ or $x = c$ ($c \in \mathbb{R}$). For example, the requirement:

The angle of joint1 shall never be greater than 170 degrees

can be rewritten as a PSP($\mathcal{D}_C$):

Globally, it is always the case that $\theta_1 < 170$

where $\theta_1 \in \mathbb{R}$ is the variable associated with the angle of *joint1* and 170 is the limiting threshold. While basic PSPs only allow for Boolean states/events in their description, PSPs($\mathcal{D}_C$) also allow for atomic numerical constraints. It is straightforward to extend the translation of [18] from basic PSPs to LTL in order to encode every PSP($\mathcal{D}_C$) to a formula in LTL($\mathcal{D}_C$). Consider, for instance, the set of requirements:

R_1 Globally, it is always the case that $\mathbf{v} \le \mathbf{5.0}$ holds.
R_2 After $\mathbf{a}$, $\mathbf{v} \le \mathbf{8.5}$ eventually holds.
R_3 After $\mathbf{a}$, it is always the case that if $\mathbf{v} \ge \mathbf{3.2}$ holds, then $\mathbf{z}$ eventually holds,

where $\mathbf{a}$ and $\mathbf{z}$ are Boolean states/events, whereas $\mathbf{v}$ is a numeric signal. These PSPs($\mathcal{D}_C$)[5] can be rewritten as the

---

following LTL($\mathcal{D}_C$) formula:

$$\begin{aligned}
&\Box(v < 5.0 \lor v = 5.0) \qquad\qquad\qquad \land \\
&\Box(a \to \Diamond(v < 8.5) \lor (v = 8.5)) \land \\
&\Box(a \to \Box(\neg(v < 3.2) \to \Diamond z)).
\end{aligned} \qquad (1)$$

Therefore, to reason about the consistency of sets of requirements written using PSPs($\mathcal{D}_C$) it is sufficient to provide an algorithm for deciding the satisfiability of LTL($\mathcal{D}_C$) formulas.

To this end, consider an LTL($\mathcal{D}_C$) formula ϕ, and let $X(\phi)$ be the set of variables and $C(\phi)$ be the set of constants that occur in ϕ. We define the *set of thresholds* $S_x(\phi) \subseteq C(\phi)$ as the set of constant values against which some variable $x \in X(\phi)$ is compared to; more precisely, for every variable $x \in X(\phi)$ we construct a set $S_x(\phi) = \{c_1, \ldots, c_n\}$ such that, for all $c_k \in \mathbb{R}$ with $1 \le k \le n$, ϕ contains a constraint of the form $x < c_k$ or $x = c_k$. For convenience, we always consider each threshold set $S_x(\phi)$ ordered in ascending order, i.e., $c_k < c_{k+1}$ for all $1 \le k < n$. For instance, in example (1), we have $X = \{v\}$ and the corresponding set of threshold is $S_v = \{3.2, 5.0, 8.5\}$. Given an LTL($\mathcal{D}_C$) formula ϕ, and some variable $x \in X(\phi)$, let $S_x(\phi) = \{c_1, \ldots, c_n\}$ be the set of thresholds for which we define the corresponding sets of *inequality propositions* $Q_x(\phi) = \{q_1, \ldots, q_n\}$ and *equality propositions* $E_x(\phi) = \{e_1, \ldots, e_n\}$. Informally, inequality propositions should be true exactly when a variable $x \in X(\phi)$ is below or between some value in the threshold set $S_x(\phi)$, whereas equality propositions should be true exactly when x is equal to some value in $S_x(\phi)$. Because of this, in our encoding we must ensure that for every computation π and time instant $i \in \mathbb{N}$ exactly one of the following cases is true ($1 \le j \le n$):

- $q_j \in \pi(i)$ for some j, $q_l \notin \pi(i)$ for all $l \ne j$ and $e_j \notin \pi(i)$ for all j;
- $e_j \in \pi(i)$ for some j, $e_l \notin \pi(i)$ for all $l \ne j$ and $q_j \notin \pi(i)$ for all j;
- $q_j \notin \pi(i)$ and $e_j \notin \pi(i)$ for all j.

The first case above corresponds to a value of x that lies between some threshold value in $S_x(\phi)$ or before its smallest value; the second case occurs when a threshold value is equal to x, and the third case is when x exceeds the highest threshold value in $S_x(\phi)$.

Given the definitions above, an LTL($\mathcal{D}_C$) formula ϕ over the set of Boolean propositions $Prop$ and the set of terms $T = C \cup X$ can be converted to an LTL formula ϕ' over the set of Boolean propositions $Prop \cup \bigcup_{x \text{ in } X}(Q_x(\phi) \cup E_x(\phi))$. We obtain this by considering, for each variable $x \in X$ and associated threshold set $S_x(\phi)$, the corresponding propositions $Q_x(\phi) = \{q_1, \ldots q_n\}$ and $E_x = \{e_1, \ldots, e_n\}$; then, for

each $c_k \in S_x(\phi)$, we perform the following substitutions:

$$x < c_k \rightsquigarrow \bigvee_{j=1}^{k} q_j \lor \bigvee_{j=1}^{k-1} e_j \quad \text{and} \quad x = c_k \rightsquigarrow e_k. \qquad (2)$$

Replacing atomic numerical constraints is not enough to ensure equisatisfiability of ϕ' with respect to ϕ. In particular, for every $x \in X(\phi)$, we must encode the informal observation made above about "mutually exclusive" Boolean valuations for propositions in $Q_x(\phi)$ and $E_x(\phi)$ as corresponding constraints:

$$\phi_M = \bigwedge_{x \in X(\phi)} \left(\bigwedge_{a,b \in M_x(\phi), a \ne b} \Box\neg(a \land b) \right), \qquad (3)$$

where $M_x(\phi) = Q_x(\phi) \cup E_x(\phi)$.

For instance, given example (1), we have $Q_v = \{q_1, q_2, q_3\}$ and $E_v = \{e_1, e_2, e_3\}$ and the mutual exclusion constraints are written as:

$$\begin{aligned}
&\phi_M = \Box\neg(q_1 \land q_2) \land \Box\neg(q_1 \land q_3) \land \Box\neg(q_1 \land e_1) \\
&\quad \land \Box\neg(q_1 \land e_2) \land \\
&\Box\neg(q_1 \land e_3) \land \Box\neg(q_2 \land q_3) \land \Box\neg(q_2 \land e_1) \\
&\quad \land \Box\neg(q_2 \land e_2) \land \\
&\Box\neg(q_2 \land e_3) \land \Box\neg(q_3 \land e_1) \land \Box\neg(q_3 \land e_2) \\
&\quad \land \Box\neg(q_3 \land e_3) \land \\
&\Box\neg(e_1 \land e_2) \land \Box\neg(e_1 \land e_3) \land \Box\neg(e_2 \land e_3).
\end{aligned} \qquad (4)$$

Therefore, the LTL formula to be tested for assessing the consistency of the requirements is

$$\begin{aligned}
\phi_M \land (\; &\Box(q_1 \lor q_2 \lor e_1 \lor e_2) \land \\
&\Box(a \to \Diamond(\textstyle\bigvee_{i=1}^{3} q_i \lor e_i)) \land \\
&\Box(a \to \Box(\neg q_1 \to \Diamond z))).
\end{aligned} \qquad (5)$$

We can now state the following:

Theorem 1 *Let ϕ be an LTL($\mathcal{D}_C$) formula on the set of proposition $Prop$ and terms $T = X(\phi) \cup C(\phi)$; for every $x \in X(\phi)$, let $S_x(\phi)$, $Q_x(\phi)$ and $E_x(\phi)$ be the corresponding set of thresholds, inequality propositions and equality propositions, respectively; let ϕ' be the LTL formula on the set of proposition $Prop \cup \bigcup_{x \in X(\phi)} Q_x(\phi) \cup E_x(\phi)$ obtained from ϕ by applying substitutions (2) for every $x \in X(\phi)$ and $c_k \in S_x(\phi)$, and let ϕ_M be the LTL formula obtained as in (3); then, the LTL($\mathcal{D}_C$) formula ϕ is satisfiable if and only if the LTL formula $\phi_M \land \phi'$ is satisfiable.*

Proof First, we prove that if ϕ is satisfiable the same holds for $\phi_M \land \phi'$. Since ϕ is satisfiable, then there exists a computation π and an evaluation ν such that $\pi, \nu \models_{\mathcal{D}_C} \phi$. Let us consider a generic variable $x \in X(\phi)$, for which the corresponding

set of thresholds is $S_x(\phi) = \{c_1, \ldots, c_n\}$. Considering that thresholds are ordered in ascending order, we construct the following sets of time instants:

$$N_{x<c_1} = \{i \in \mathbb{N} \mid \pi, \nu, i \models_{\mathcal{D}_C} x < c_1\}$$
$$N_{x=c_1} = \{i \in \mathbb{N} \mid \pi, \nu, i \models_{\mathcal{D}_C} x = c_1\}$$
$$N_{c_1<x<c_2} = \{i \in \mathbb{N} \mid \pi, \nu, i \models_{\mathcal{D}_C} x > c_1 \wedge x < c_2\}$$
$$\cdots$$
$$N_{c_{n-1}<x<c_n} = \{i \in \mathbb{N} \mid \pi, \nu, i \models_{\mathcal{D}_C} x > c_{n-1} \wedge x < c_n\}$$
$$N_{x=c_n} = \{i \in \mathbb{N} \mid \pi, \nu, i \models_{\mathcal{D}_C} x = c_n\}$$
$$N_{x>c_n} = \{i \in \mathbb{N} \mid \pi, \nu, i \models_{\mathcal{D}_C} x > c_n\}$$

which, given the standard semantics of "<" and "=", are a partition of $\mathbb{N}$. Let $\mathcal{N}_x$ denote such partition for a specific variable $x \in X(\phi)$. We construct a computation π' such that, for all time instants $i \in \mathbb{N}$ and propositions $p \in Prop$, we have $p \in \pi'(i)$ exactly when $p \in \pi(i)$ and, for each variable $x \in X(\phi)$, given $Q_x(\phi) = \{q_1, \ldots q_n\}$ and $E_x(\phi) = \{e_1, \ldots e_n\}$, we have also

- $q_1 \in \pi'(i)$ exactly when $i \in N_{x<c_1}$;
- $e_1 \in \pi'(i)$ exactly when $i \in N_{x=c_1}$;
- $q_2 \in \pi'(i)$ exactly when $i \in N_{c_1<x<c_2}$;
- $\ldots$
- $q_n \in \pi'(i)$ exactly when $i \in N_{c_{n-1}<x<c_n}$;
- $e_n \in \pi'(i)$ exactly when $i \in N_{x=c_n}$.

Notice that for all $i \in N_{x>c_n}$, we have that $\pi'(i) \cap M_x(\phi) = \emptyset$, where $M_x(\phi) = Q_x(\phi) \cup E_x(\phi)$. Since $\mathcal{N}_x$ is a partition of $\mathbb{N}$ for each variable $x \in X(\phi)$, it follows that $\pi' \models \phi_M$ because for all $a, b \in M_x(\phi)$, there is no time instant $i \in \mathbb{N}$ such that $\pi', i \models a \wedge b$. Now, we show that for every $i \in \mathbb{N}$, $\pi', i \models \phi'$ if and only if $\pi, \nu, i \models_{\mathcal{D}_C} \phi$ by induction on the set of subformulas of ϕ. Let ψ and ψ' be two subformulas of ϕ and ϕ', respectively. For every $i \in \mathbb{N}$:

- if $\psi \equiv p$ for $p \in Prop$ then $\psi' \equiv p$; therefore, for any given $i \in \mathbb{N}$, we have $\pi, \nu, i \models_{\mathcal{D}_C} p$ if and only if $\pi', i \models p$ by construction of π'.
- if $\psi \equiv (x < c_k)$ for some variable $x \in X(\phi)$ and some constant $c_k \in S_x(\phi)$ then, according to (2),

$$\psi' \equiv \bigvee_{j=1}^{k} q_j \vee \bigvee_{j=1}^{k-1} e_j.$$

Let $N_{x,k}$ be the set defined as

$$N_{x,k} = N_{x<c_1} \cup N_{x=c_1} \cup \ldots \cup N_{c_{k-1}<x<c_k}.$$

There are two cases: either $i \in N_{x,k}$ or $i \notin N_{x,k}$. In the former case, we have that $\pi, \nu, i \models_{\mathcal{D}_C} (x < c_k)$ and, by construction of π', this happens exactly when $\pi', i \models q_j$

for some $1 \leq j \leq k$ or $\pi', i \models e_j$ for some $1 \leq j < k$ which, by the semantics of disjunction and construction of π', is also exactly when $\pi', i \models \psi'$. In the second case, $\pi, \nu, i \not\models_{\mathcal{D}_C} (x < c_k)$ and, by construction of π', this happens exactly when $\pi', i \not\models q_j$ for all $1 \leq j \leq k$ and $\pi', i \not\models e_j$ for all $1 \leq j < k$ which, by the semantics of disjunction, is also exactly when $\pi', i \not\models \psi'$.

- if $\psi \equiv x = c_k$ for some variable $x \in X(\phi)$ and some constant $c_k \in S_x(\phi)$ then, according to (2), $\psi' \equiv e_k$. The time instants $i \in \mathbb{N}$ in which $\pi, \nu, i \models_{\mathcal{D}_C} x = c_k$ are contained in the set $N_{x=c_k}$, so there are two cases: either $i \in N_{x=c_k}$ or $i \notin N_{x=c_k}$. In the former case, we have that $\pi, \nu, i \models_{\mathcal{D}_C} (x = c_k)$ and, by construction of π', this happens exactly when $\pi', i \models e_k$. In the second case, $\pi, \nu, i \not\models_{\mathcal{D}_C} (x = c_k)$ and, by construction of π', this happens exactly when $\pi', i \not\models e_k$.

- if $\psi \equiv \neg\alpha$ then $\psi' = \neg\alpha'$; by induction, we can assume that for every i, we have $\pi, \nu, i \models_{\mathcal{D}_C} \alpha$ if and only if $\pi', i \models \alpha'$, and thus $\pi, i \not\models_{\mathcal{D}_C} \alpha$ if and only if $\pi', i \not\models \alpha'$. By the semantics of negation, we have that for any given $i \in \mathbb{N}, \pi, i, \nu \models_{\mathcal{D}_C} \neg\alpha$ if and only if $\pi, i, \nu \not\models_{\mathcal{D}_C} \alpha$ and this happens exactly when $\pi', i \not\models \alpha'$, i.e., $\pi', i \models \neg\alpha'$;

- if $\psi \equiv (\alpha \vee \beta)$ then $\psi' \equiv \alpha' \vee \beta'$; by induction, we can assume that for all $i \in \mathbb{N}$ we have that $\pi, \nu, i \models_{\mathcal{D}_C} \alpha$ and $\pi, \nu, i \models_{\mathcal{D}_C} \beta$ if and only if $\pi', i \models \alpha'$ and $\pi', i \models \beta'$, respectively. By the semantics of disjunction, we have that, for any given $i \in \mathbb{N}, \pi, i, \nu \models_{\mathcal{D}_C} \alpha \vee \beta$ exactly when $\pi, \nu, i \models \alpha$ or $\pi, \nu, i \models \beta'$ and this happens exactly when $\pi', i \models \alpha'$ or $\pi', i \models \beta'$, i.e., by the semantics of disjunction, $\pi', i \models \alpha' \vee \beta'$.

- if $\psi \equiv \mathcal{X} \alpha$ then $\psi' \equiv \mathcal{X} \alpha'$; by induction, we can assume that for all $j \in \mathbb{N}$ we have $\pi, \nu, j \models_{\mathcal{D}_C} \alpha$ if and only if $\pi', j \models \alpha'$. By the semantics of the "next" operator, we have that, for any given $i \in \mathbb{N}, \pi, i, \nu \models_{\mathcal{D}_C} \mathcal{X} \alpha$ if and only if $\pi, \nu, i + 1 \models_{\mathcal{D}_C} \alpha$ which happens exactly when $\pi', i + 1 \models \alpha'$, i.e., $\pi', i \models \mathcal{X} \alpha$.

- if $\psi \equiv \alpha \, \mathcal{U} \, \beta$ then $\psi' = \alpha' \, \mathcal{U} \, \beta'$; by induction, we can assume that, for all $j \in \mathbb{N}$, we have $\pi, \nu, j \models_{\mathcal{D}_C} \beta$ if and only if $\pi', j \models \beta'$ and that $\pi, \nu, j \models_{\mathcal{D}_C} \alpha$ if and only if $\pi', j \models \alpha'$. By the semantics of the "until" operator, we have that, for any given $i, \pi, i, \nu \models_{\mathcal{D}_C} \alpha \, \mathcal{U} \, \beta$ if and only if for some $j \geq i$ we have $\pi, \nu, j \models_{\mathcal{D}_C} \beta$ and for all k such that $i \leq k < j$ we have $\pi, \nu, k \models_{\mathcal{D}_C} \alpha$. However, the former happens exactly when for the same $j \in \mathbb{N}$ we have $\pi', j \models \beta'$ and for all k such that $i \leq k < j$ we have $\pi', k \models_{\mathcal{D}_C} \alpha'$, i.e., $\pi', i \models \alpha' \, \mathcal{U} \, \beta'$.

We now prove that the satisfiability of $\phi_M \wedge \phi'$ in LTL implies the satisfiability of ϕ in LTL$(\mathcal{D}_C)$. First, we observe that, for a generic variable $x \in X(\phi)$, and for all time instants $i \in \mathbb{N}$, every computation π' such that $\pi' \models \phi_M$ has at most one proposition $p \in M_x(\phi)$ for which $p \in \pi(i)$. Therefore, for all variables $x \in X(\phi)$ and for every time instant $i \in \mathbb{N}$,

we have the following cases only (where $n = |S_x(\phi)| = |E_x(\phi)| = |Q_x(\phi)|$):

1. $\pi', i \models e_k$ for some $e_k \in E_x(\phi)$; consequently, as long as $k < n$, also $\pi', i \models \bigvee_{j=1}^{k+1} q_j \vee \bigvee_{j=1}^{k} e_j$ holds.
2. $\pi', i \models q_k$ for some $q_k \in Q_x(\phi)$, and thus $\pi', i \models \bigvee_{j=1}^{k} q_j \vee \bigvee_{j=1}^{k-1} e_j$ holds.
3. $\pi', i \not\models p$ for every $p \in M_x(\phi)$; consequently, for all k it is also the case that $\pi', i \not\models \bigvee_{j=1}^{k} q_j \vee \bigvee_{j=1}^{k-1} e_j$ and $\pi', i \not\models e_k$.

A computation π and an evaluation ν such that $\pi, \nu \models_{\mathcal{D}_C} \phi$ can be constructed as follows. For every $p \in Prop$, and time instant $i \in \mathbb{N}$, let $p \in \pi(i)$ exactly when $p \in \pi'(i)$. As for the evaluation ν, for a generic variable $x \in X(\phi)$, and for every time instant $i \in \mathbb{N}$, we can construct ν considering that π' is bound to satisfy the three cases above :

1. $\nu(x, i) = c_k$ for the same k s.t. $\pi', i \models e_k$; consequently, as long as $k < n$, both $\pi, \nu, i \models x < c_{k+1}$ and $\pi, \nu, i \models x = c_k$ hold.
2. $\nu(x, i) = v$ and, for the same k s.t. $\pi', i \models q_k$, if $k > 1$, then $c_{k-1} < v < c_k$, else if $k = 1$, then $v < c_1$; consequently $\pi, \nu, i \models x < c_k$ holds and, in case $k > 1$, $\pi, \nu, i \not\models x < c_j$ for all $j < k$.
3. $\nu(x, i) = v$ with $v > c_n$; consequently $\pi, \nu, i \not\models x < c_k$ and $\pi, \nu, i \not\models x = c_k$ for all k

An induction proof analogous to the one provided for the "if" part can be provided to show that if $\pi' \models \phi'$, then also $\pi, \nu \models \phi$, with π and ν constructed as shown above. $\qquad\square$

The proposed translation from LTL($\mathcal{D}_C$) to a LTL formula is also quite compact, i.e., the number of symbols in the LTL encoding grows at most quadratically with the number of symbols in the original formula. Let us define the size of a formula ϕ, denoted as $|\phi|$, in the usual way, i.e., by counting the number of symbols in it. We can state the following:

Theorem 2 *Let ϕ be an LTL($\mathcal{D}_C$) formula on the set of proposition $Prop$ and terms $T = X(\phi) \cup C(\phi)$; for every $x \in X(\phi)$, let $S_x(\phi)$, $Q_x(\phi)$ and $E_x(\phi)$ be the corresponding set of thresholds, inequality propositions and equality propositions, respectively; let ϕ' be the LTL formula on the set of proposition $Prop \cup \bigcup_{x \in X(\phi)} Q_x(\phi) \cup E_x(\phi)$ obtained from ϕ by applying substitutions (2) for every $x \in X(\phi)$ and $c_k \in S_x(\phi)$, and ϕ_M be the LTL formula obtained as in (3); the size of $\phi' \wedge \phi_M$ is at most quadratic in the size of ϕ, i.e., $O(|\phi' \wedge \phi_M|) = O(|\phi|^2)$.*

Proof From Eq. (3), for each variable $x \in X(\phi)$, all combinations of two elements from the set $M_x(\phi) = Q_x(\phi) \cup E_x(\phi)$

are required to build ϕ_M. Therefore, if $n = |S_x(\phi)|$, the number of conjuncts of the form $\Box \neg(a \wedge b)$ in ϕ_M is

$$\binom{2n}{2} = \frac{2n!}{2!(2n-2)!} = \frac{2n(2n-1)}{2} = n(2n-1). \qquad (6)$$

If we consider $m = \max_{x \in X(\phi)} |S_x(\phi)|$ and the number of conjuncts derived in Eq. (6), it follows that

$$|\phi_M| = O(|X(\phi)| \cdot m(2m-1)) = O(|X(\phi)| \cdot m^2). \qquad (7)$$

Now, it remains to show the effect of substitution (2) in ϕ. For every variable $x \in X(\phi)$ and for each constant $c_k \in S_x(\phi)$ in ϕ, we have:

- one proposition in ϕ' for each occurrence of the term $x = c_k$ in ϕ;
- a formula of size $2k - 1$ in ϕ' for each occurrence of the term $x < c_k$ in ϕ.

Let $m = \max_{x \in X(\phi)} |S_x(\phi)|$, and p be the maximum number of occurrences in ϕ of any condition $x = c$ or $x < c$ for specific values of $x \in X(\phi)$ and $c \in C(\phi)$. Then, we can write

$$|\phi| = O(|X(\phi)| \cdot p \cdot m + r), \qquad (8)$$

where r is the number of symbols that are not terms. Since each term in ϕ is translated to a formula of size $O(m)$ in ϕ', we have that

$$|\phi'| = O(|X(\phi)| \cdot p \cdot m^2 + r). \qquad (9)$$

Considering (7) together with (9), we obtain

$$O(|\phi' + \phi_M|) = O(|X(\phi)| \cdot m^2) + O(|X(\phi)| \cdot p \cdot m^2 + r)$$
$$= O(|X(\phi)| \cdot m^2 \cdot (1 + p) + r). \qquad (10)$$

Given (8) and the fact that the values of the parameters $|X(\phi)|$, p and r do not depend on the translation, from (10) we conclude that $O(|\phi' + \phi_M|) = O(|\phi|^2)$. $\qquad\square$

4 Inconsistency explanation

Given a set $R = \{r_1, \ldots, r_n\}$ of inconsistent requirements written as PSP($\mathcal{D}_C$), the aim of the algorithms proposed in this section is to compute a minimal unsatisfiable core (MUC), i.e., a subset $I \subseteq R$ such that removing any element r_i from I makes the set consistent again. Table 1 shows an inconsistent specification as a set $R = \{r_1, \ldots r_7\}$ of seven requirements. Looking at the table, we can see that there are 4 different MUCs in R, namely $\{r_1, r_2\}$, $\{r_2, r_6\}$, $\{r_3, r_4, r_5\}$,

Table 1 Set R of inconsistent PSPs

r_i	PSP
r_1	Globally, it is always the case that A holds
r_2	Globally, it is never the case that A holds
r_3	Globally, it is always the case that B holds
r_4	Globally, it is always the case that if B holds, then C holds as well
r_5	Globally, it is never the case that C holds
r_6	Globally, it is always the case that A and B holds
r_7	After B, D eventually holds

$\{r_4, r_5, r_6\}$. In the remainder of the section, we present two algorithms devoted to the extraction of MUC for PSPs.

4.1 Linear deletion-based MUC extraction

The first algorithm we present is based on a deletion-based strategy, and its pseudo-code is depicted in Algorithm 1. The procedure works as follows. If the set $R' \leftarrow R \setminus \{r\}$ with $r \in R$ is inconsistent, then r is not in the MUC. On the other hand, if R' is consistent, then r is part of a MUC and cannot be removed. Such operation is repeated iteratively and the algorithm terminates when all requirements have been checked for inclusion in the MUC.

Algorithm 1 Linear Deletion-Based MUC Extraction Algorithm

```
1: function FINDINCONSISTENCY(R)
2:     R' ← R
3:     for r_i ∈ R do
4:         R' ← R' \ {r_i}
5:         if ISCONSISTENT(R') then
6:             R' ← R' ∪ {r_i}
7:         end if
8:     end for
9:     return R'
10: end function
```

It is easy to see that, with $|R| = n$, the loop iterates n times, and that at each iteration the ISCONSISTENT function is called once. The input of the function is R' and its size is given by $|R'|$. The number of elements in R' is reduced by one at each iteration, but r_i could be added back again in R', depending on the result of ISCONSISTENT. The worst case is obtained when all requirements are part of the MUC, i.e., each requirement r_i is first removed and then reinserted again. In this case, the model checker is called each time with $n - 1$ requirements. The overall complexity is therefore $O(n \cdot C(n))$, where n is the number of elements initially in R and $C(n)$ is the complexity for the consistency check of n requirements. The algorithm is therefore linear in the number of calls to the model checker.

Example 1 Considering the set R in Table 1, Algorithm 1 works along the following steps. The final result is $R' =$

Step	r_i	R'	ISCONSISTENT(R')
1:	r_1	$\{r_2, r_3, r_4, r_5, r_6, r_7\}$	FALSE
2:	r_2	$\{r_3, r_4, r_5, r_6, r_7\}$	FALSE
3:	r_3	$\{r_4, r_5, r_6, r_7\}$	FALSE
4:	r_4	$\{r_5, r_6, r_7\}$	TRUE
5:	r_5	$\{r_4, r_6, r_7\}$	TRUE
6:	r_6	$\{r_4, r_5, r_7\}$	TRUE
7:	r_7	$\{r_4, r_5, r_6\}$	FALSE

$\{r_4, r_5, r_6\}$. It is worth to notice that this result depends on the extraction order of the requirements. It is easy to see that processing the requirements in reverse order would yield $R' = \{r_1, r_2\}$ as a result instead.

4.2 Dichotomic MUC extraction

Algorithm 2 is based on the same general-purpose structure of Algorithm 1, but it also exploits the fact that the dimension of the MUC is often much smaller than $|R|$. Therefore, it is possible to exploit a "divide and conquer" strategy to reduce the search space. Considering Algorithm 2, R is split in two halves R_1 and R_2, such that $R_1 \cup R_2 = R$ and $R_1 \cap R_2 = \emptyset$. If one of the two halves (plus I) is inconsistent, then there is no need to explore the other one and we can proceed recursively. Otherwise, it means that the MUC has been split into two halves and further search is needed. This is done by means of two recursive calls (lines 21–22); the former performs the search on R_2 considering the whole set R_1 as inconsistent, while the latter continues the search on R_1, removing from I the requirements that still need to be checked. The algorithm terminates when R has 1 or 0 elements.

As for the complexity of the algorithm the best case occurs when the MUC is always in the first half of R. In such a case, half of the requirements are discarded at each iteration, and it is easy to see that complexity is $\Omega(\log |R|)$. The worst case occurs when the set of inconsistent requirements I coincides with R. Taking into account Table 1, let R be comprised of $\{r_1, r_2, r_3, r_4\}$ and let MUC be R itself. At the first step, the algorithm checks $R'_1 = \{r_1, r_2\}$ and $R'_2 = \{r_3, r_4\}$, but both sets are consistent. Therefore, FINDINCONSISTENCY is called recursively with $R = \{r_3, r_4\}$ and $I = \{r_1, r_2\}$. At this point, we have $R''_1 = \{r_3\}$ and $R''_2 = \{r_4\}$. The algorithm checks the consistency of $\{r_1, r_2, r_3\}$ and $\{r_1, r_2, r_4\}$ and returns to the previous recursive call. This time FINDINCONSISNTENCY is called again, but with $R = \{r_1, r_2\}$ and $I = \{r_3, r_4\}$ and the same process is applied. In general, if $|R| = n$ and $C(n)$ is the complexity for the consistency check of n requirements, then the worst case complexity of this algorithm is $O(n \cdot C(n))$ – the same as the previous one. However, as we will show

Algorithm 2 Dichotomic MUC Extraction Algorithm

```
 1: function FINDINCONSISTENCY(R)
 2:     return FINDINCONSISTENCY(R, ∅)
 3: end function

 4: function FINDINCONSISTENCY(R, I)
 5:     if |R| ≤ 1 then
 6:         if ISCONSISTENT(I) then
 7:             return I ∪ R
 8:         else
 9:             return I
10:         end if
11:     end if
12:     (R₁, R₂) ← SPLIT(R)
13:     if |R₁| > 1 and |R₂| > 1 then
14:         if ¬ ISCONSISTENT(R₁ ∪ I) then
15:             return FINDINCONSISTENCY(R₁, I)
16:         end if
17:         if ¬ ISCONSISTENT(R₂ ∪ I) then
18:             return FINDINCONSISTENCY(R₂, I)
19:         end if
20:     end if
21:     I ← FINDINCONSISTENCY(R₂, I ∪ R₁)
22:     I ← FINDINCONSISTENCY(R₁, I \ R₁)
23:     return I
24: end function
```

in Sect. 5.2, when $|I| \ll |R|$ it is noticeable faster than the linear version.

Example 2 Considering again the set R reported in Table 1, in the following we report step-by-step how Algorithm 2 works. For lack of space in the table, we replace ISCONSISTENT with $isCons$.

Step	R	R_1	R_2	I	$isCons(R_1 \cup I)$	$isCons(R_2 \cup I)$
1:	$\{r_1, \ldots, r_7\}$	$\{r_1, r_2, r_3\}$	$\{r_4, r_5, r_6, r_7\}$	$\{\}$	$False$	$-$
2:	$\{r_1, r_2, r_3\}$	$\{r_1\}$	$\{r_2, r_3\}$	$\{\}$	$True$	$True$
3:	$\{r_2, r_3\}$	$\{r_2\}$	$\{r_3\}$	$\{r_1\}$	$-$	$-$
4:	$\{r_2\}$	$-$	$-$	$\{r_1\}$	$-$	$-$
5:	$\{r_1\}$	$-$	$-$	$\{r_2\}$	$-$	$-$

In the first step, the algorithm splits the initial set R into two subsets R_1 and R_2 and checks the consistency of the first one. Since R_1 is inconsistent, the algorithm automatically discards R_2 and continue with step 2. Also, in this case the new set $R = \{r_1, r_2, r_3\}$ is split into two, but this time both are consistent and so the two recursive calls in line 21–22 are executed: the first one is resolved in steps 3 and 4, while the second one in step 5. In the last two steps, the basic case is reached (lines 5–11), and since the call to ISCONSISTENT(I) returns true in both cases, r_1 and r_2 are added to I. Therefore, $I = \{r_1, r_2\}$ is returned as final answer. In this case, ISCONSISTENT is called 6 times instead of 7 as in the previous example.

5 Analysis with probabilistic requirement generation

The aim of this section is twofold. On the one hand, we evaluate the scalability of our approach for consistency checking, experimenting the encoding proposed in Sect. 3 with a pool of state-of-the-art LTL model checkers. On the other hand, we assess the performance of the MUC extraction algorithms described in Sect. 4, in order to evaluate the possibility of their usage in contexts of practical interest.

Since we want to have control over different dimensions of the specifications—the kind of requirements, the number of constraints and the size of the corresponding domains—we generate artificial specifications using a probabilistic model that we devised and implemented specifically to carry out the experiments herein presented.

In particular, the following parameters can be tuned in our generator of specifications:

- The number of requirements generated ($\#req$).
- The probability of each different body to occur in a pattern.
- The probability of each different scope to occur in a pattern.
- The size ($\#vars$) of the set from which variables are picked uniformly at random to build patterns.
- The size (dom) of the domain from which the thresholds of the atomic constraints are chosen uniformly at random.

5.1 Evaluation of LTL($\mathcal{D}_c$) satisfiability

The goal of this experiment is to evaluate the performance—in terms of correctness, efficiency and scalability—of LTL model checkers for the consistency checking task described in Sect. 3. To this end, we evaluate the performances of state-of-the-art tools for LTL satisfiability, and then we consider the best among such tools to assess whether our approach can scale to sets of requirements of realistic size. All the experiments here reported ran on a workstation equipped with 2 Intel Xeon E5-2640 v4 CPUs and 256 GB RAM running Debian with kernel 3.16.0-4.

Evaluation of LTL satisfiability solvers The tools considered in our analysis are the ones included in the portfolio solver POLSAT [25], namely AALTA [28], NUSMV [9], PLTL [46] and TRP++ [21]. We also consider LEVIATHAN [7], a tableau-based system for consistency checking that has been recently published. Notice that in the case of NUSMV, we consider two different encodings. With reference to Property 1, the first encoding defines ϕ_M as an invariant—denoted as NUSMV- INVAR—and ϕ' is the property to check; the second encoding considers $\phi_M \wedge \phi$ as the property to check—denoted as NUSMV- NOINVAR. Finally, concerning AALTA, we slightly modified its default version in order to be able

Table 2 Evaluation of LTL satisfiability solvers on randomly generated requirements

dom	2				4				8				16			
#vars	16		32		16		32		16		32		16		32	
Tool	**S**	**T**	**S**	**T**	**S**	**T**	**S**	**T**	**S**	**T**	**S**	**T**	**S**	**T**	**S**	**T**
AALTA	**16**	0.0	**27**	0.1	**22**	0.1	**29**	0.4	**26**	0.6	**29**	1.4	**25**	2.8	**31**	4.9
LEVIATHAN	4	0.1	6	0.3	7	0.8	5	0.2	0	–	7	2.3	4	47.7	7	12.8
NuSMV- INVAR	11	30.4	10	185.1	10	804.2	9	881.3	11	68.1	8	402.9	10	1172.6	8	1001.9
NuSMV- NOINVAR	11	65.0	10	489.7	7	303.6	7	505.5	11	92.4	10	1277.6	8	660.0	9	1394.5
PLTL	8	25.0	11	108.1	9	1.2	10	0.6	10	19.6	11	0.1	11	14.5	14	3.5

The first line reports the size of the domain (*dom*), while the second line reports the total amount of variables (*vars*) for each domain size. Then, for each tool (on the first column), the table shows the total amount of solved problems and the CPU time (in s) spent to solve them (columns "S" and "T", respectively)

to evaluate large formulas. In particular, we modified the source code increasing of two orders of magnitude the input size buffer.

In our experimental analysis, we set the range of the parameters as follows: $\#vars \in \{16, 32\}, dom \in \{2, 4, 8, 16\}$ and $\#req \in \{8, 16, 32, 64\}$. For each combination of the parameters with $v \in \#vars$, $r \in \#req$ and $d \in dom$, we generate 10 different benchmarks. Each benchmark is a specification containing r requirements where each scope has (uniform) probability 0.2 and each body has (uniform) probability 0.1. Then, for each atomic numerical constraint in the benchmark, we choose a variable out of v possible ones, and a threshold value out of d possible ones. In Table 2, we show the results of the analysis. Notice that we do not show the results of TRP++ because of the high number of failures obtained. Looking at the table, we can see that AALTA is the tool with the best performances, as it is capable of solving two times the problems solved by other solvers in most cases. Moreover, AALTA is up to 3 orders of magnitude faster than its competitors. Considering unsolved instances, it is worth noticing that in our experiments AALTA never reaches the granted time limit (10 CPU min), but it always fails beforehand. This is probably due to the fact that AALTA is still in a relatively early stage of development and it is not as mature as NuSMV and PLTL. Most importantly, we did not find any discrepancies in the satisfiability results of the evaluated tools, with the noticeable exception of TRP++, for which we did not report performance in Table 2.

Evaluation of scalability The analysis involves 2560 different benchmarks generated in the previous experiment. The initial value of $\#req$ has been set to 15, and it has been doubled until 1920, thus obtaining benchmarks with a total amount of requirements equal to 15, 30, 60, 120, 240, 480, 960 and 1920. Similarly has been done for $\#vars$ and $\#dom$; the former ranges from 5 to 640, while the latter ranges from 4 to 32. At the end of the generation, we obtained 10 different sets composed of 256 benchmarks. In Figs. 2 and 3, we present the results, obtained running AALTA. The figure is composed

of 8 plots, one for each value of $\#vars$. Looking at the plots in Figs. 2 and 3, we can see that the difficulty of the problem increases when all the values of the considered parameters increase, and this is particularly true considering the total amount of requirements. The parameter $\#dom$ has a higher impact of difficulty when the number of variables is small. Indeed, when the number of variables is less then 40 there is a clear difference between solving time with $\#dom = 4$ and $\#dom = 32$. On the other hand, when the number of variables increases, all the plots for various values of $\#dom$ are very close to each other. As a final remark, we can see that even considering the largest problem ($\#vars = 640$, $\#dom = 32$), more than 60% of the problems are solved by AALTA within the time limit of 10 min.

5.2 Evaluation of MUC extraction

In order to evaluate the algorithms proposed in Sect. 4, we consider the pool of inconsistent benchmarks resulting from the experiment presented in Sect. 5.1, for a total amount of 559, having different requirements set dimension as reported in Table 3. All the experiments reported in this section ran on a workstation equipped with an Intel Xeon E31245 @ 3.30 GHz CPU and 16 GB RAM running Ubuntu 14.04 LTS.

In Fig. 4, we report the results obtained from the experiment described above. For each plot, we report the median CPU time (in s) over 10 runs of the same benchmark, granting for each run 600 CPU s. AALTA has been used for the satisfiability check.

Looking at the plots, we can see that the dichotomic algorithm is, as expected, overall faster than the linear one. Despite the fact that they show similar performance for benchmarks having 8 and 16 requirements (top-most plots in Fig. 4), looking at the plots in the middle of Fig. 4 we can see that the dichotomic algorithm is at least one order of magnitude faster than the linear one for benchmarks having 32 and 60 requirements. More, we report that the latter was able to return MUCs only for 62 out of 65 and 43 out
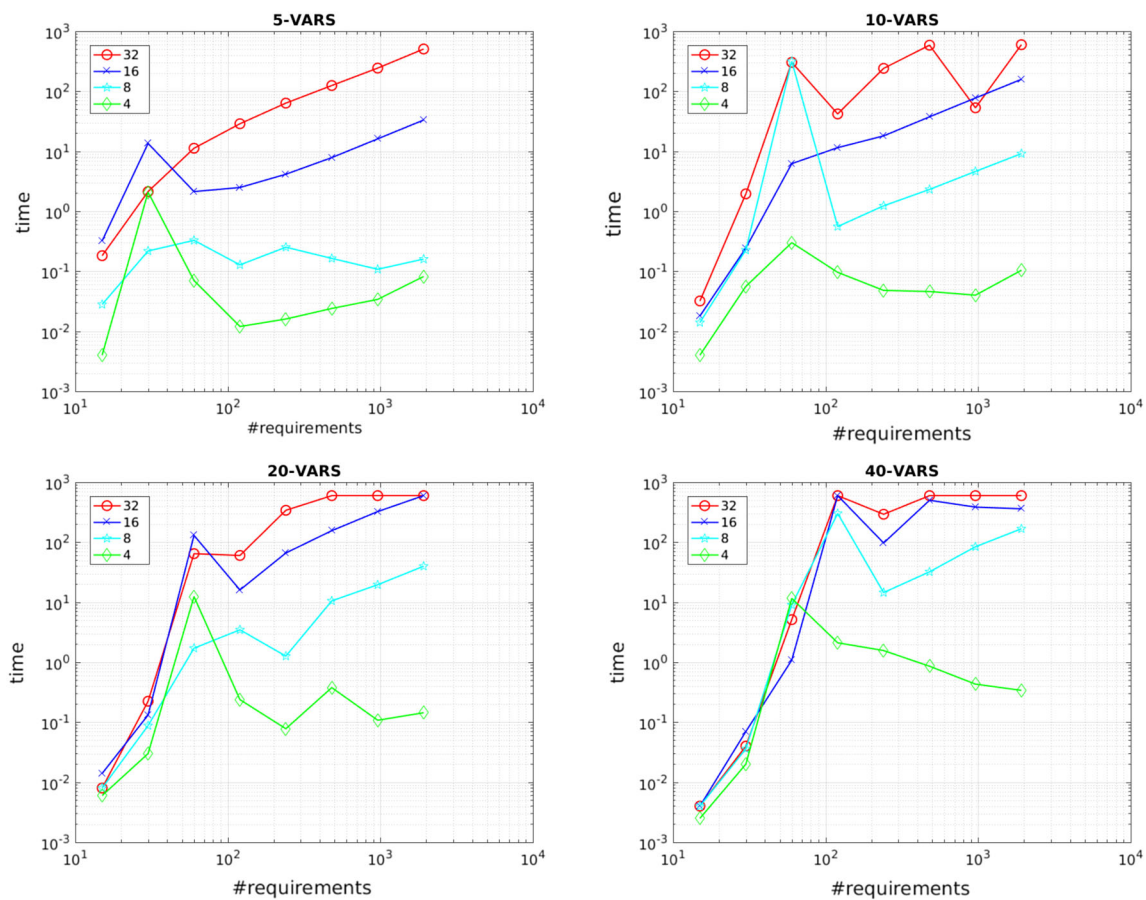
Fig. 2 Scalability analysis (Part 1). On the *x*-axes (*y*-axes resp.), we report *req* (CPU time in s resp.). Both the axes are in logarithmic scale. In each plot, we consider different values of #*dom*. In particular, the diamond green line is for #*dom* = 4, the light blue line with stars is for #*dom* = 8, the blue crossed lines and red circled ones denote #*dom* = 16 and #*dom* = 32, respectively (color figure online)

of 83, while the former returned a solution for all instances with 32 requirements and 81 out of 83 for instances with 60 requirements.

Considering the plots in the bottom of Fig. 4, we can see that the gap between the two algorithms increases even further: the linear one was able to return MUCs only for 34 and 12 benchmarks of 120 and 240 requirements, respectively, while the dichotomic one returned a MUC for 138 out of 147 and 168 out of 210 benchmarks. In addition, it is worth noticing that the MUCs found are usually small in size; indeed, in all 6 configurations, the median size of the MUCs found by the two algorithms is 2.

Finally, we report that we involved in our analysis also benchmarks composed of 480 requirements, but our algorithms were not able to return a solution within the considered CPU time limit.

As a final remark, notice that we limit the presentation of the results to the algorithms presented in Sect. 4 because state-of-the-art tools able to cope with this task, namely PLTL- MUP [20] and TRP++UC [45], report the same correct-ness and scalability issues of their counterparts presented in Sect. 5.1. For instance, considering the benchmark with 60 requirements—the first one for which we can see a notice-able difference between the performance of the linear and the dichotomic algorithm—we report that PLTL- MUP was not able to solve any instance, while TRP++UC tops its perfor-mance at 37% of our worst algorithm (the linear one solved 43 instances out of 83). We also involved in our preliminary analysis also PROCMINE [1], but we do not report its results for similar motivations.

6 Analysis with a controller for a robotic manipulator

In this section, as a basis for our experimental analysis, we consider a set of requirements from the design of an embedded controller for a robotic manipulator. The con-troller should direct a properly initialized robotic arm—and related vision system—to look for an object placed in a
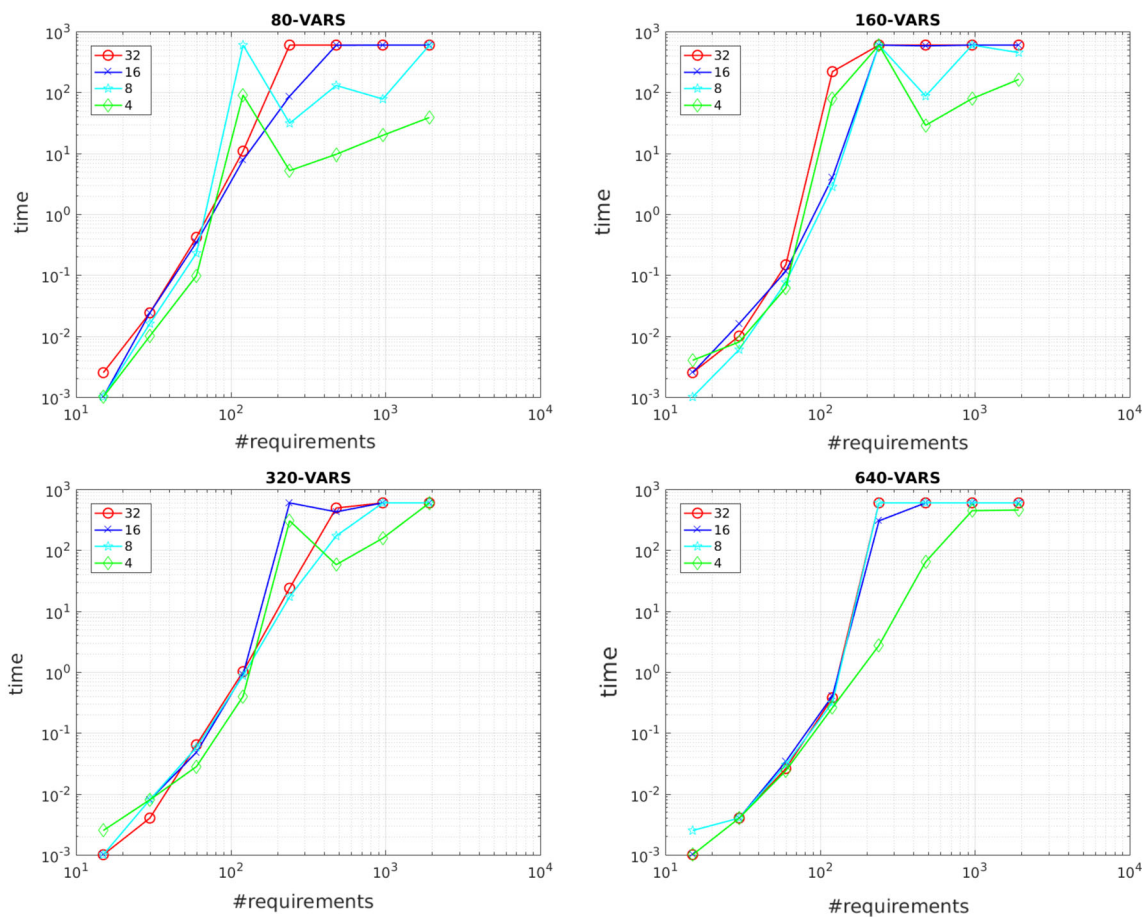
Fig. 3 Scalability analysis (Part 2). Plots are organized as in Fig. 2

Table 3 Synopsis of the pool of benchmarks involved in the analysis of MUC extraction algorithms

#req	N
8	16
16	38
32	65
60	83
120	147
240	210

The table is organized in two columns, namely the total amount of requirements for each benchmark (column "#req") and the total amount of benchmarks falling in the related category (column "N")

given position and move to such position in order to grab the object; once grabbed, the object is to be moved into a bucket placed in a given position and released without touching the bucket. The robot must stop also in the case of an unintended collision with other objects or with the robot itself — collisions can be detected using torque estimation from current

sensors placed in the joints. Finally, if a general alarm is detected, e.g., by the interaction with a human supervisor, the robot must stop as soon as possible. The manipulator is a four-degree-of-freedom Trossen Robotics WidowX arm[6] equipped with a gripper: Fig. 5 shows a snapshot of the robot in the intended usage scenario taken from V-REP[7] simulator. The design of the embedded controller is currently part of the activities related to the "Self-Healing System for Planetary Exploration" use case [35] in the context of the EU project CERBERO.

In this case study, constrained numerical signals are used to represent requirements related to various parameters, namely angle, speed, acceleration and torque of the 4 joints, size of the object picked, and force exerted by the end effector. We consider 75 requirements, including those involving scenario-independent constraints like joints limits, and mutual exclusion among states, as well as specific requirements related to the conditions to be met at each state. The set of requirements involved in our analysis includes 14

---

[6] http://www.trossenrobotics.com/widowxrobotarm.

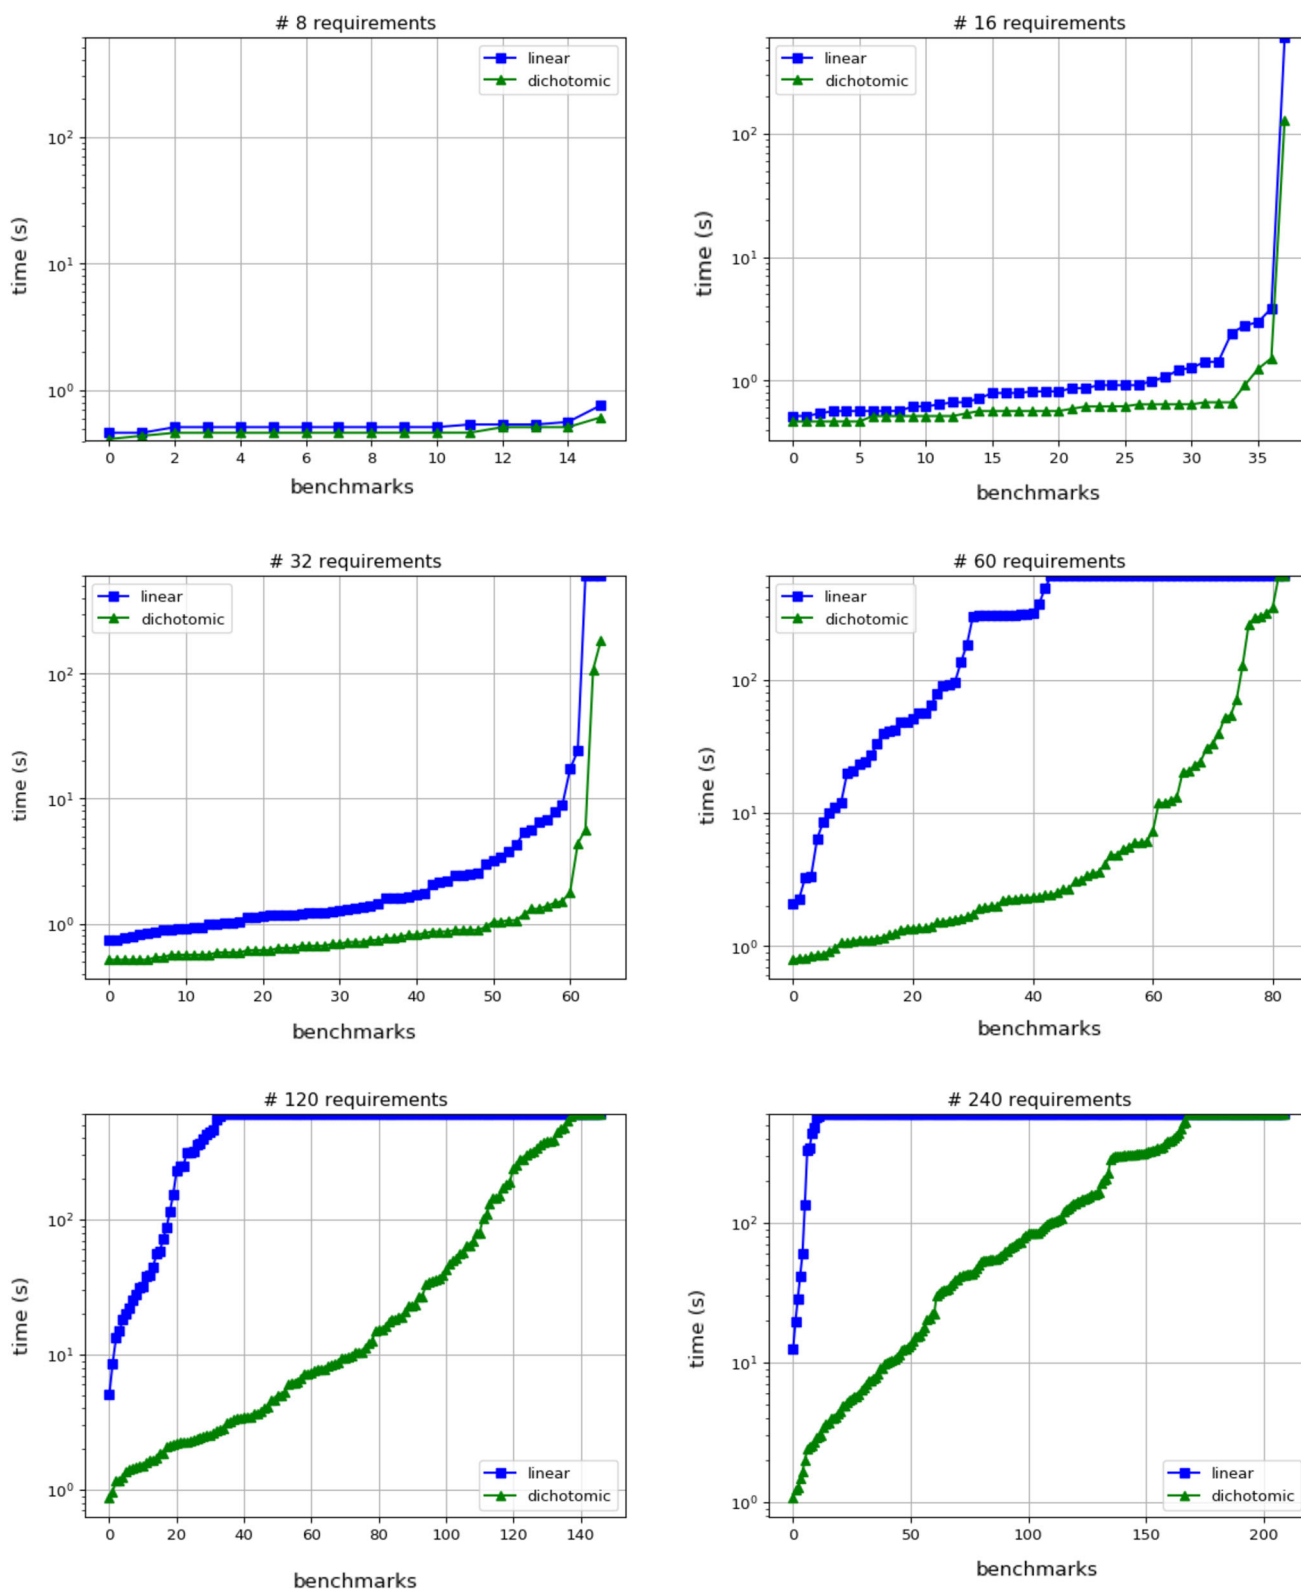[7] http://www.coppeliarobotics.com/.

Fig. 4 Performance of the algorithms for MUC extraction. On the x-axes we report the number of benchmarks, and on the y-axes we report the time in logarithmic scale. In each plot, we consider different values of #req. The green and blue lines show median times of the dichotomic and linear algorithms, respectively (color figure online)

**Fig. 5** WidowX robotic arm (left) and the simulated arm moving a grabbed object in the bucket on the left (right)

Table 4 Robotic use case requirements synopsis

Pattern	Specification			Fault injections		
	AFTER	AFTER_UNTIL	GLOBALLY	AFTER	AFTER_UNTIL	GLOBALLY
Absence	–	12	14	[F4]	–	[F3]
Existence	9	–	–	–	[F5]	[F4, F6]
Invariant	–	–	29	–	–	[F2, F6]
Precedence	–	–	1	–	–	–
ResponseChain	–	–	2	–	–	–
Response	1	–	4	–	–	[F1]
Universality	2	–	1	–	–	–

The table is organized as follows: the first column reports the name of the patterns, and it is followed by two groups of three columns denoted with the scope type: the first group refers to the intended specification, and the second to the one with fault injections. Each cell in the first group reports the number of requirements grouped by pattern and by scope type. Cells in the second group categorize the 6 injected faults, labeled with F1, …, F6

Boolean signals and 20 numerical ones. In Table 4, we present a synopsis of the requirements, to give an idea of the kind of patterns used in the specification.[8] While most requirements are expressed with the invariant pattern, e.g., mutual exclusiveness of states and safety conditions, the expressivity of LTL is required to describe the evolution of the system. Indeed, as shown in [18] and [41], it is often the case that few PSPs cover the majority of specifications, whereas others are sparsely used.

Our first experiment[9] is to run NuSMV- INVAR on the intended specification translated to LTL($\mathcal{D}_C$). The motivation for presenting the results with NuSMV- INVAR rather than AALTA is twofold: while its performances are worse than AALTA, NuSMV- INVAR is more robust in the sense that it either reaches the time limit or it solves the problem, without ever failing for unspecified reasons like AALTA does at times; second, it turns out that NuSMV- INVAR can deal flawlessly and in reasonable CPU times with all the specifications we consider in this section, both the intended one and the ones obtained by injecting faults. In particular, on the intended

specification, NuSMV- INVAR is able to find a valid model for the specification in 37.1 CPU s, meaning that there exists at least a model able to satisfy all the requirements simultaneously. Notice that the translation time from patterns to formulas in LTL($\mathcal{D}_C$) is negligible with respect to the solving time. Our second experiment is to run NuSMV- INVAR on the specification with some faults injected. In particular, we consider six different faults, and we extend the specification in six different ways considering one fault at a time. The patterns related to the faults are summarized in Table 4. In case of faulty specifications, NuSMV- INVAR concludes that there is no model able to satisfy all the requirements simultaneously. In particular, in the case of F2 and F3, NuSMV- INVAR returned the result in 2.1 and 1.7 CPU s, respectively. Concerning the other faults, the tools were one order of magnitude slower in returning the satisfiability result. In particular, it spent 16.8, 50.4, 12.2 and 25.6 CPU s in the evaluation of the requirements when faults 1, 4, 5 and 6 are injected, respectively.

The noticeable difference in performances when checking for different faults in the specification is mainly due to the fact that F2 and F3 introduce an initial inconsistency, i.e., it would not be possible to initialize the system if they were present in the specification, whereas the remaining faults introduce inconsistencies related to interplay among con-

---

[8] The full list of requirements and the fault injection examples are available at https://github.com/SAGE-Lab/robot-arm-usecase.

[9] Experiments herein presented ran on a PC equipped with a CPU Intel Core i7-2760QM @ 2.40 GHz (8 cores) and 8 GB of RAM, running Ubuntu 14.04 LTS.

straints in time, and thus additional search is needed to spot problems. In order to explain this difference, let us first consider fault 2:

> *Globally, it is always the case that if* `state_init` *holds, then not* `arm_idle` *holds as well.*

It turns out that in the intended specification there is one requirement specifying exactly the opposite, i.e., that when the robot is in `state_init`, then `arm_idle` must hold as well. Thus, the only models that satisfy both requirements are the ones preventing the robot arm to be in `state_init`. However, this is not possible because other requirements related to the state evolution of the system impose that `state_init` will eventually occur and, in particular, that it should be the first one. On the other hand, if we consider fault 6:

> *Globally, it is always the case that if* `arm_moving` *holds, then* `joint1_speed` > *15.5 holds as well.*
> *Globally,* `arm_moving` *and* `proximity_sensor` *= 10.0 eventually holds.*

we can see that the first requirement sets a lower speed bound at 15.5 deg/s for `joint1` when the arm is moving, while there exists a requirement in the intended specification setting an upper speed bound at 10 deg/s when the proximity sensor detects an object closer than 20 cm. In this case, the model checker is still able to find a valid model in which `proximity_sensor` < 20.0 never happens when `arm_moving` holds, but the second requirement in fault 6 prohibits this opportunity. It is exactly this kind of interplay among different temporal properties which makes NuSMV- INVAR slower in assessing the (in)consistency of some specifications.

7 Conclusions

In this paper, we have extended basic PSPs over the constraint system $\mathcal{D}_C$, and we have provided an encoding from any PSP($\mathcal{D}_C$) into a corresponding LTL formula. This enables us to deal with the satisfiability of specifications of practical interest and to verify them using state-of-the-art reasoning tools currently available for LTL. Noticeably, even considering the largest problem in our experiments (#$vars$ = 640, #dom = 32), more than 60% of the problems are solved (by AALTA) within the time limit of 10 min. Overall, using the specifications generated with our probabilistic model we have shown that our approach implemented on the tool AALTA scales to problems containing more than a thousand requirements over hundreds of variables. Considering a real-world case study in the context of the EU project CERBERO, we have shown that it is feasible to check specifications and uncover injected faults, even with tools other than AALTA.

Inconsistency explanations could be provided for all, but the largest specifications in our benchmark base. These results witness that our approach is viable and worth of adoption in the process of requirement engineering. Our next steps toward this goal will include easing the translation from natural language requirements to patterns and extending the pattern language to deal with other relevant aspects of cyberphysical systems, such as real-time constraints and related logics (e.g., signal temporal logic [32]). Further elements will also include domain-specific strategies to search for MUCs in requirements aiming at improving the performance of the algorithm presented in Sect. 4, i.e., discovering or approximating the minimum set of requirements causing the inconsistency while looking for the consistency of the set, instead to do it at the end of the consistency checking, as we did in Sect. 4.

References

1. Awad A, Goré R, Thomson J, Weidlich M (2011) An iterative approach for business process template synthesis from compliance rules. In: International conference on advanced information systems engineering, Springer, pp 406–421
2. Bakker RR, Dikker F, Tempelman F, Wognum PM (1993) Diagnosing and solving over-determined constraint satisfaction problems. IJCAI 93:276–281
3. Barnat J, Bauch P, Beneš N, Brim L, Beran J, Kratochvíla T (2016) Analysing sanity of requirements for avionics systems. Form Asp Comput 28(1):45–63
4. Belov A, Marques-Silva J (2011) Accelerating MUS extraction with recursive model rotation. In: Formal methods in computer-aided design (FMCAD), 2011. pp 37–40. IEEE
5. Belov A, Marques-Silva J (2012) Muser2: an efficient mus extractor. J Satisf Boolean Model Comput 8:123–128
6. Bendík J (2017) Consistency checking in requirements analysis. In: Proceedings of the 26th ACM SIGSOFT international symposium on software testing and analysis, pp 408–411. ACM
7. Bertello M, Gigante N, Montanari A (2016) Reynolds M Leviathan: A new ltl satisfiability checking tool based on a one-pass tree-shaped tableau. In: Proceedings of the twenty-fifth international joint conference on artificial intelligence, AAAI Press, pp 950–956. IJCAI'16, http://dl.acm.org/citation.cfm?id=3060621.3060753
8. Chinneck JW, Dravnieks EW (1991) Locating minimal infeasible constraint sets in linear programs. ORSA J Comput 3(2):157–168
9. Cimatti A, Clarke E, Giunchiglia E, Giunchiglia F, Pistore M, Roveri M, Sebastiani R, Tacchella A (2002) NuSMV 2: an opensource tool for symbolic model checking. In: 14th international conference on computer aided verification (CAV 2002), pp 359–364
10. Clarke EM, Emerson EA, Sistla AP (1986) Automatic verification of finite-state concurrent systems using temporal logic specifications. ACM Trans Progr Lang Syst (TOPLAS) 8(2):244–263
11. Comon H, Cortier V (2000) Flatness is not a weakness. In: International workshop on computer science logic, pp 262–276

12. Demri S, DSouza D (2007) An automata-theoretic approach to constraint LTL. Inf Comput 205(3):380–415

13. Desrosiers C, Galinier P, Hertz A, Paroz S (2009) Using heuristics to find minimal unsatisfiable subformulas in satisfiability problems. J Comb Optim 18(2):124–150

14. Dillon LK, Kutty G, Moser LE, Melliar-Smith PM, Ramakrishna YS (1994) A graphical interval logic for specifying concurrent systems. ACM Trans Softw Eng Methodol (TOSEM) 3(2):131–165

15. Dokhanchi A, Hoxha B, Fainekos G (2015) Metric interval temporal logic specification elicitation and debugging. In: 13th ACM-IEEE international conference on formal methods and models for codesign, pp 21–23

16. Dokhanchi A, Hoxha B, Fainekos G (2016) Formal requirement debugging for testing and verification of cyber-physical systems. arXiv preprint arXiv:1607.02549

17. Dravnieks EW (1989) Identifying minimal sets of inconsistent constraints in linear programs: deletion, squeeze and sensitivity filtering. Carleton University

18. Dwyer MB, Avrunin GS, Corbett JC (1999) Patterns in property specifications for finite-state verification. In: Proceedings of the 21st international conference on software engineering, pp 411–420

19. Fisman D, Kupferman O, Sheinvald-Faragy S, Vardi MY (2008) A framework for inherent vacuity. In: Haifa verification conference, Springer, pp 7–22

20. Goré R, Huang J, Sergeant T, Thomson J (2013) Finding minimal unsatisfiable subsets in linear temporal logic using BDDS. https://www.timsergeant.com/files/pltlmup/gore_huang_sergeant_thomson_mus_pltl.pdf

21. Hustadt U, Konev B (2003) TRP++ 2.0: A temporal resolution prover. In: 19th international conference on automated deduction, pp 274–278

22. Junker U (2001) Quickxplain: Conflict detection for arbitrary constraint propagation algorithms. In: IJCAI01 workshop on modelling and solving problems with constraints

23. Kesten Y, Manna Z, McGuire H, Pnueli A (1993) A decision algorithm for full propositional temporal logic. In: International conference on computer aided verification, Springer, pp 97–109

24. Konrad S, Cheng BH (2005) Real-time specification patterns. In: Proceedings of the 27th international conference on software engineering, pp 372–381

25. Li J, Pu G, Zhang L, Yao Y, Vardi M Y et al. (2013) Polsat: A portfolio LTL satisfiability solver. arXiv preprint arXiv:1311.1602

26. Li J, Yao Y, Pu G, Zhang L, He J (2014) Aalta: an LTL satisfiability checker over infinite/finite traces. In: Proceedings of the 22nd ACM SIGSOFT international symposium on foundations of software engineering, pp 731–734

27. Li J, Zhang L, Pu G, Vardi MY, He J (2013) LTL satisfiability checking revisited. In: 20th international symposium on temporal representation and reasoning, pp 91–98

28. Li J, Zhu S, Pu G, Vardi MY (2015) Sat-based explicit LTL reasoning. In: 11th Haifa verification conference, pp 209–224

29. Liffiton MH, Malik A (2013) Enumerating infeasibility: finding multiple muses quickly. In: International conference on AI and OR techniques in constraint programming for combinatorial optimization problems, Springer, pp 160–175

30. Liffiton MH, Sakallah KA (2008) Algorithms for computing minimal unsatisfiable subsets of constraints. J Autom Reason 40(1):1–33

31. Lumpe M, Meedeniya I, Grunske L (2011) PSPWizard: machine-assisted definition of temporal logical properties with specification patterns. In: Proceedings of the 19th ACM SIGSOFT symposium and the 13th European conference on foundations of software engineering, pp 468–471

32. Maler O, Nickovic D (2004) Monitoring temporal properties of continuous signals. In: Formal techniques, modelling and analysis of timed and fault-tolerant systems, Springer, pp 152–166

33. Manna Z, Pnueli A (2012) Temporal verification of reactive systems: safety. Springer, Berlin

34. Marques-Silva J, Lynce I (2011) On improving MUS extraction algorithms. In: International conference on theory and applications of satisfiability testing. Springer, pp 159–173

35. Masin M, Palumbo F, Myrhaug H, de Oliveira Filho J, Pastena M, Pelcat M, Raffo L, Regazzoni F, Sanchez A, Toffetti A, et al. (2017) Cross-layer design of reconfigurable cyber-physical systems. In: 2017 design, automation & test in Europe conference & exhibition (DATE), pp 740–745. IEEE

36. Nadel A (2010) Boosting minimal unsatisfiable core extraction. In: Proceedings of the 2010 conference on formal methods in computer-aided design, pp 221–229. FMCAD Inc

37. Narizzano M, Pulina L, Tacchella A, Vuotto S (2018) Consistency of property specification patterns with boolean and constrained numerical signals. In: NASA formal methods: 10th international symposium, NFM 2018, Newport News, VA, USA, April 17–19, 2018, Proceedings, vol 10811. Springer, pp 383–398

38. Pnueli A (1977) The temporal logic of programs. In: 18th annual symposium on foundations of computer science, 1977, pp 46–57. IEEE

39. Pnueli A, Manna Z (1992) The temporal logic of reactive and concurrent systems. Springer, Berlin 16:12

40. Pnueli A, Rosner R (1989) On the synthesis of a reactive module. In: Proceedings of the 16th ACM SIGPLAN-SIGACT symposium on principles of programming languages, pp 179–190. ACM

41. Post A, Hoenicke J (2012) Formalization and analysis of real-time requirements: a feasibility study at BOSCH. Verified software: theories, tools, experiments, pp 225–240

42. Rozier KY, Vardi MY (2007) LTL satisfiability checking. In: Spin. vol. 4595, pp 149–167. Springer

43. Rozier KY, Vardi MY (2010) LTL satisfiability checking. Int J Softw Tools Technol Transf (STTT) 12(2):123–137

44. Rozier KY, Vardi MY (2011) A multi-encoding approach for LTL symbolic satisfiability checking. In: International symposium on formal methods, Springer, pp 417–431

45. Schuppan V (2016) Extracting unsatisfiable cores for lTl via temporal resolution. Acta Inform 53(3):247–299

46. Schwendimann S (1998) A new one-pass tableau calculus for PLTL. In: International conference on automated reasoning with analytic tableaux and related methods, Springer, pp 277–291

47. Wolper P (1985) The tableau method for temporal logic: an overview. Logique et Analyse 28(110/111):119–136