

Towards an Effective Decision Procedure for LTL formulas with Constraints*

Marco Comini¹, Laura Titolo¹, and Alicia Villanueva²

¹ DIMI, Università degli Studi di Udine,
{marco.comini,laura.titolo}@uniud.it

² DSIC, Universitat Politècnica de València
villanue@dsic.upv.es

Abstract. This paper presents an ongoing work that is part of a more wide-ranging project whose final scope is to define a method to validate LTL formulas w.r.t. a program written in the timed concurrent constraint language *tccp*, which is a logic concurrent constraint language based on the concurrent constraint paradigm of Saraswat. Some inherent notions to *tccp* processes are non-determinism, dealing with partial information in states and the monotonic evolution of the information.

In order to check an LTL property for a process, our approach is based on the abstract diagnosis technique. The concluding step of this technique needs to check the validity of an LTL formula (with constraints) in an effective way.

In this paper, we present a decision method for the validity of temporal logic formulas (with constraints) built by our abstract diagnosis technique.

1 Introduction

The *ccp* paradigm is different from other (concurrent) programming paradigms mainly due to the notion of store-as-constraint that replaces the classical store-as-valuation model. It is based on an underlying constraint system that handles constraints on variables, thus, it deals with partial information. One challenging characteristics of the *ccp* framework is that programs can manifest non-monotonic behaviors, implying that standard approaches cannot be directly adapted. Within this family, [7] introduced the *Timed Concurrent Constraint Language (tccp)* by adding to the original *ccp* model the notion of time and the ability to capture the absence of information. With these features, one can specify behaviors typical of reactive systems such as *timeouts* or *preemption* actions.

Modeling and verifying concurrent systems by hand can be very complicated. Thus, the development of automatic formal methods is essential. One of the most known technique for formal verification is model checking, that was originally

* This work has been partially supported by the EU (FEDER) and the Spanish MEC/MICINN, ref. TIN 2010-21062-C02-02, and by Generalitat Valenciana, ref. PROMETEO2011/052.

introduced in [3,16] to automatically check if a finite-state system satisfies a given property. It consisted in an exhaustive analysis of the state-space of the system; thus the state-explosion problem is its main drawback and, for this reason, many proposals in the literature try to mitigate it.

All the proposals of model checking have in common that a subset of the model of the (target) program has to be built, and sometimes the needed fragment is quite huge. Our final goal is to define a method that avoids the need to build the model of a system in order to check the validity of some temporal property. We propose an extension of the abstract diagnosis technique of [5] so that the abstract domain is formed by LTL formulas. The final step of the method consists in checking whether a given formula, built from the program (abstract) semantics and the specification, is valid.

In this work, we present a decision procedure to check the validity of such formulas. The linear temporal logic used in this work is an adaptation of the propositional LTL logic to the concurrent constraint framework, following the ideas of [15,8,9,18]. It is expressive enough to represent the abstract semantics of *tccp* with much precision and we provide a decision procedure based on the tableaux method of [11,13]. As we show through this paper, the considered logic has some differences w.r.t. the classic LTL logic due to the constraint nature and to the fact that models for *tccp* programs have some special characteristics (inherited from the *ccp* paradigm).

2 The small-step denotational semantics of *tccp*

The *tccp* language [7] is particularly suitable to specify both reactive and time critical systems. As the other languages of the *ccp* paradigm [17], it is parametric w.r.t. a cylindric constraint system which handles the data information of the program in terms of constraints. The computation progresses as the concurrent and asynchronous activity of several agents that can (monotonically) accumulate information in a *store*, or query some information from it. Briefly, a cylindric constraint system³ is an algebraic structure $\mathbf{C} = \langle \mathcal{C}, \leq, \otimes, true, false, Var, \exists \rangle$ composed of a set of constraints \mathcal{C} such that (\mathcal{C}, \leq) is a complete algebraic lattice where \otimes is the *lub* operator and *false* and *true* are respectively the greatest and the least element of \mathcal{C} ; *Var* is a denumerable set of variables and \exists existentially quantifies variables over constraints. The *entailment* \vdash is the inverse of order \leq .

Given a cylindric constraint system \mathbf{C} and a set of process symbols Π , the syntax of agents is given by the grammar:

$$A ::= \text{skip} \mid \text{tell}(c) \mid A \parallel A \mid \exists x A \mid \sum_{i=1}^n \text{ask}(c_i) \rightarrow A \mid \text{now } c \text{ then } A \text{ else } A \mid p(\vec{x})$$

where c, c_1, \dots, c_n are finite constraints in \mathbf{C} ; $p_{/m} \in \Pi$, $x \in Var$ and $\vec{x} \in Var \times \dots \times Var$. A *tccp* program P is an object of the form $D.A$, where A is an agent, called *initial agent*, and D is a set of *process declarations* of the form $p(\vec{x}) :- A$.

³ See [7,17] for more details on cylindric constraint systems.

The notion of time is introduced by defining a discrete and global clock. The *ask*, *tell* and *process call* agents take one time-unit to be executed.

Intuitively, the *skip* agent represents the successful termination of the agent computation. The *tell*(c) agent adds the constraint c to the current store and stops. It takes one time-unit, thus the constraint c is visible to other agents from the following time instant. The store is updated by means of the \otimes operator of the constraint system. The choice agent $\sum_{i=1}^n \text{ask}(c_i) \rightarrow A_i$ consults the store and non-deterministically executes (at the following time instant) one of the agents A_i whose corresponding guard c_i holds in the current store; otherwise, if no guard is entailed by the store, the agent suspends. The agent *now* c *then* A_1 *else* A_2 behaves in the current time instant like A_1 (respectively A_2) if c is (respectively is not) entailed by the store. Entailment is checked by using the \vdash operator of the constraint system. Note that this agent can process negative information: it can capture when some information is not present in the store since the agent A_2 is executed both when $\neg c$ is entailed, but also when neither c nor $\neg c$ are entailed. $A_1 \parallel A_2$ models the parallel composition of A_1 and A_2 in terms of maximal parallelism, i.e., all the enabled agents of A_1 and A_2 are executed at the same time. The agent $\exists x A$ is used to make variable x local to A . To this end, it uses the $\tilde{\exists}$ operator of the constraint system. Finally, the agent $p(\bar{x})$ takes from D a declaration of the form $p(\bar{x}) :- A$ and executes A at the following time instant. We assume that the set D of declarations is closed w.r.t. parameter names.

In this work, we refer to the denotational concrete semantics defined in [5] which is fully-abstract w.r.t. the small-step behavior of *tccp*. Due to space limitations, we just introduce intuitively the most relevant aspects of such semantics and provide a global view of it by means of an example. The missing definitions, as well as the proofs of all the results, can be found in [6].

The denotational semantics of a *tccp* program consists of a set of *conditional (timed) traces* that represent, in a compact way, all the possible behaviors that the program can manifest when fed with an *input* (initial store). The set of all conditional timed traces is denoted with \mathbf{M} . Intuitively, conditional traces can be seen as hypothetical computations in which, for each time instant, a condition represents the information that the global store has to satisfy in order to proceed to the next time instant. Briefly, a conditional trace is a (possibly infinite) sequence $t_1 \cdots t_n \cdots$ of *conditional states*, which can be of three forms:

- conditional store:** a pair $\eta \succ c$, where η is a *condition* and $c \in \mathbf{C}$ a store;
- stuttering:** the construct $\text{stutt}(C)$, with $C \subseteq \mathbf{C} \setminus \{\text{true}\}$;
- end of a process:** the construct \boxtimes .

The conditional store $\eta \succ c$ is used to represent a hypothetical computation step, where η is the condition that the current store must satisfy in order to make the computation proceed. c represents the information that is provided by the program up to the current time instant. A *condition* η is a pair $\eta = (\eta^+, \eta^-)$ where $\eta^+ \in \mathbf{C}$ and $\eta^- \in \wp(\mathbf{C})$ are called positive and negative condition, respectively. The positive/negative condition represents information that a given store must/must not entail, thus they have to be consistent in the sense that $\forall c^- \in \eta^-, \eta^+ \not\# c^-$.

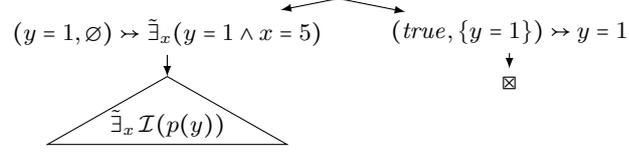


Fig. 1. Tree representation of $\mathcal{A}[[A]]_{\mathcal{I}}$ of Example 1.

Conditional states in a conditional trace have to be monotone (i.e., for each $t_i = \eta_i \succcurlyeq c_i$ and $t_j = \eta_j \succcurlyeq c_j$ such that $j \geq i$, $c_j \vdash c_i$) and consistent (i.e., each store must not entail its associated negative conditions).

We abuse in notation and define as $\exists_x r$ the sequence resulting by removing from r all the information about the variable x . Moreover, $r \in \mathbf{M}$ is said to be *self-sufficient* if the first condition is $(true, \emptyset)$ and, for each $t_i = (\eta_i^+, \eta_i^-) \succcurlyeq c_i$ and $t_{i+1} = (\eta_{i+1}^+, \eta_{i+1}^-) \succcurlyeq c_{i+1}$, $c_i \vdash \eta_{i+1}^+$ (each store satisfies the successive condition). Moreover, r is *self-sufficient w.r.t. $x \in \mathcal{V}$* (*x -self-sufficient*) if $\exists_{\mathcal{V}ar \setminus \{x\}} r$ is self-sufficient. Intuitively, for self-sufficient conditional traces, no additional information (from other agents) is needed in order to *complete* the computation.

Example 1. Let us consider a program with a single process declaration $D := \{p(y) :- A\}$, where

$$A := \exists x (\text{now } y = 1 \text{ then } (\text{tell}(x = 5) \parallel p(y)) \text{ else } \text{tell}(y = 1))$$

The semantics of A is graphically represented in Fig. 1. The left branch represents the computation when the information $y = 1$ is entailed by the store (the positive condition requires $y = 1$), thus the information added by the tell agent ($x = 5$) is joined to the store and, in the following state, since it is invoked a (recursive) process call, we find the interpretation of the process called (represented by the triangle labeled with the interpretation \mathcal{I}). Process calls do not modify the store when invoked, but they affect the store from the following time instant. The right branch is taken only if $y = 1$ does not hold in the current state, then it holds at the following time instant due to the tell agent.

3 Constraint System Linear Temporal Logic

In this section, we define a variation of the classical Linear Temporal Logic [14]. Following [15,8,9,18], the idea is to replace atomic propositions by constraints of the underlying constraint system. This logic is the basis for the definition of the abstract semantics needed in our abstract diagnosis technique.

Definition 1 (csLTL formulas). *Given a cylindric constraint system \mathbf{C} , $c \in \mathbf{C}$ and $x \in \text{Var}$, formulas of the csLTL Logic over \mathbf{C} are defined as:*

$$\phi ::= true \mid false \mid c \mid \dot{\neg} \phi \mid \phi \dot{\wedge} \phi \mid \dot{\exists}_x \phi \mid \bigcirc \phi \mid \phi \mathcal{U} \phi.$$

We denote with csLTL the set of all temporal formulas over \mathbf{C} .

The formulas *true*, *false*, $\dot{\neg}\phi$, and $\phi_1 \wedge \phi_2$ have the classical logical meaning. The atomic formula $c \in \mathbf{C}$ states that c has to be entailed by the current store. $\dot{\exists}_x \phi$ is the existential quantification over the set of variables Var . $\bigcirc \phi$ states that ϕ holds at the next time instant, while $\phi_1 \mathcal{U} \phi_2$ states that ϕ_2 eventually holds and in all previous instants ϕ_1 holds. In the sequel, we use $\phi_1 \dot{\vee} \phi_2$ as a shorthand for $\dot{\neg}(\dot{\neg}\phi_1 \wedge \dot{\neg}\phi_2)$; $\phi_1 \dot{\rightarrow} \phi_2$ for $\dot{\neg}\phi_1 \dot{\vee} \phi_2$; $\diamond \phi$ for *true* $\mathcal{U} \phi$ and $\square \phi$ for $\dot{\neg}\diamond\dot{\neg}\phi$. A *constraint formula* is an atomic formulas c or its negation $\dot{\neg}c$. Formulas of the form $\bigcirc \phi$ and $\dot{\neg}\bigcirc \phi$ are called *next* formulas. Constraint and next formulas are said to be *elementary* formulas. Finally, formulas of the form $\phi_1 \mathcal{U} \phi_2$ (or $\diamond \phi$ or $\dot{\neg}(\square \phi)$) are called *eventualities*.

The truth of a formula $\phi \in \mathbf{csLTL}$ is defined w.r.t. a trace $r \in \mathbf{M}$. As usually done in the context of temporal logics, we define the satisfaction relation \models only for infinite conditional traces. We implicitly transform finite traces (which end in \boxtimes) by replicating the last store infinite times.

Definition 2. The semantics of $\phi \in \mathbb{F}$ is given by $\gamma^{\mathbb{F}}: \mathbb{F} \rightarrow \mathbb{M}$ defined as

$$\gamma^{\mathbb{F}}(\phi) := \bigsqcup \{r \in \mathbf{M} \mid r \models \phi\}, \quad (3.1)$$

where, for each $\phi, \phi_1, \phi_2 \in \mathbf{csLTL}$, $c \in \mathbf{C}$ and $r \in \mathbf{M}$, the satisfaction relation \models is defined as:

$$r \models \textit{true} \quad (3.2a)$$

$$r \not\models \textit{false} \quad (3.2b)$$

$$(\eta^+, \eta^-) \rightsquigarrow d \cdot r' \models c \quad \textit{iff} \quad \eta^+ \vdash c \quad (3.2c)$$

$$\textit{stutt}(\eta^-) \cdot r' \models c \quad \textit{iff} \quad \forall d^- \in \eta^-. c \not\vdash d^- \textit{ and } r' \models c \quad (3.2d)$$

$$r \models \dot{\neg}\phi \quad \textit{iff} \quad r \not\models \phi \quad (3.2e)$$

$$r \models \phi_1 \wedge \phi_2 \quad \textit{iff} \quad r \models \phi_1 \textit{ and } r \models \phi_2 \quad (3.2f)$$

$$r \models \dot{\exists}_x \phi \quad \textit{iff} \quad \exists r' \textit{ s.t. } \dot{\exists}_x r' = \dot{\exists}_x r, r' \textit{ x-self-sufficient, } r' \models \phi \quad (3.2g)$$

$$t \cdot r \models \bigcirc \phi \quad \textit{iff} \quad r \models \phi \quad (3.2h)$$

$$r \models \phi_1 \mathcal{U} \phi_2 \quad \textit{iff} \quad \exists i \geq 1. \forall j < i. r^i \models \phi_2 \textit{ and } r^j \models \phi_1 \quad (3.2i)$$

We abuse of notation and we extend \models to sets of formulas: $r \models \Phi \iff \forall \phi \in \Phi. r \models \phi$. A formula ϕ is said to be *satisfiable* if there exists $r \in \mathbf{M}$ such that $r \models \phi$, while it is said to be *valid* if, for all $r \in \mathbf{M}$, $r \models \phi$.

Example 2. The formula $\square(x > 0 \dot{\rightarrow} \bigcirc z > 1)$ expresses that forever, whenever $x > 0$ is entailed by the store, then $z > 1$ is entailed at the following time instant.

4 Abstract diagnosis of temporal properties

Abstract diagnosis is a semantic based method to identify bugs in programs. It was originally defined for logic programming [4] and then extended to other paradigms [1,2,10,5]. This technique is based on the definition of an abstract semantics for the program which must be a sound approximation of its behavior.

Then, given an abstract specification \mathcal{S} of the expected behavior of the program, the abstract diagnosis technique automatically detects the errors in the program by checking if the result of one computation of the semantics evaluation function (where the procedure calls are interpreted over \mathcal{S}) is “contained” in the specification itself.

A first approach to the abstract diagnosis of *tccp* was presented in [5] by using as specifications sets of abstract traces. The main drawback of that proposal was that specifications were given in terms of traces, thus they can be tedious to write. In order to overcome that problem, we have defined an abstract semantics in terms of csLTL formulas. This allows us to express the intended behavior in a more compact way, by means of a csLTL formula.

The semantics evaluation function $\mathcal{A}[[A]]$, given an agent A and an interpretation \mathcal{I} (for the process symbols of A), builds a csLTL formula representing a correct approximation of the small-step behavior of A . The semantics of the declarations is given in terms of the fixpoint of a semantics evaluation function $\mathcal{D}[[D]]$ which associates to each procedure declaration $p(\vec{x})$ the logic disjunction of the abstract semantics of every agent A such that $p(\vec{x}) :- A$ belongs to the declaration D (i.e., $\mathcal{D}[[D]]_{\mathcal{I}}(p(\vec{x})) := \bigvee_{p(\vec{x}) :- A \in D} \mathcal{A}[[A]]_{\mathcal{I}}$).

The *abstract diagnosis* technique determines exactly the “originating” symptoms and, in the case of incorrectness, the faulty process declaration in the program. This is captured by the definitions of *abstractly incorrect process declaration* and *abstract uncovered element*. Informally, a process declaration D is abstractly incorrect if it derives a wrong abstract element $\phi_t \in \text{csLTL}$ from the intended semantics \mathcal{S} . Dually, ϕ_t is uncovered if the declarations cannot derive it from the intended semantics.

We show here the main result of abstract diagnosis, which determines the form of formulas that we need to check for validity.

Theorem 1. *Consider a set of declarations D and an abstract specification \mathcal{S} .*

1. *If there are no abstractly incorrect process declarations in D (i.e., $\mathcal{D}[[D]]_{\mathcal{S}} \dot{\rightarrow} \mathcal{S}$), then D is partially correct w.r.t. \mathcal{S} (the small-step denotational semantics of D is “contained” in the semantics of \mathcal{S}).*
2. *Let D be partially correct w.r.t. \mathcal{S} . If D has abstract uncovered elements then D is not complete.*

Therefore, in order to check partial correctness of a program, it is sufficient to check the implication $\mathcal{D}[[D]]_{\mathcal{S}} \dot{\rightarrow} \mathcal{S}$.

Because of the approximation, it can happen that a (concretely) correct declaration is abstractly incorrect. Hence, abstract incorrect declarations are in general just a warning about a possible source of errors. However, an abstract correct declaration cannot contain an error; therefore, no (manual) inspection is needed for declarations which are not signalled. Moreover, it happens that all concrete errors—that are “visible”—are detected, as they lead to an abstract incorrectness or abstract uncovered.

Example 3. Assume we need to check that the program in Example 1 satisfies that the constraint $y = 1$ is eventually entailed by the store. The intended specification for the process p corresponding to this property is $\mathcal{S}(p(y)) := \diamond(y = 1)$. The csLTL-semantic \mathcal{D} for $p(y)$ with the given specification as interpretation is

$$\mathcal{D}[\![D]\!]_{\mathcal{S}}(p(y)) = \exists_x ((y = 1 \wedge \bigcirc x = 5 \wedge \bigcirc(\diamond y = 1)) \dot{\vee} (\neg y = 1 \wedge \bigcirc y = 1))$$

Note that the resulting formula has a clear correspondence with the behavior of the program. We have two disjuncts, one for each branch of the conditional in the body of the declaration. Moreover, since the conditional agent is in the scope of a local variable x , both disjuncts are enclosed within an existential quantification. The computation of the csLTL-semantic is compositional, based on the structure of the program. The first disjunct corresponds to the case when the guard of the conditional agent is satisfied ($y = 1$), thus in the following time instant two things happen: $x = 5$ is entailed, and also is entailed the interpretation for the process call p because a recursive call is run. This illustrates how the intended specification is used as the interpretation of process calls.

Following Theorem 1, to check whether the process $p(y)$ satisfies the property, it is sufficient to show that $\mathcal{D}[\![D]\!]_{\mathcal{S}}(p(y)) \dot{\rightarrow} \mathcal{S}(p(y))$ is valid.

The fact that our technique is based on abstract diagnosis becomes explicit when two situations occur simultaneously. When the process that is being analyzed has more than one fixpoint (this essentially happens when D contains a loop which does not produce contributes at all) and the specification is an eventuality that is not inconsistent with the program behavior, then it may happen that the actual behaviour does not model a specification \mathcal{S} , which is a non-least fixpoint of $\mathcal{D}[\![D]\!]$, but we do not detect abstractly incorrect declarations since \mathcal{S} is a fixpoint. This is natural in the context of the abstract diagnosis technique: what we are proving is that, if $\mathcal{S}(p(\vec{x}))$ is assumed to hold for each process $p(\vec{x})$ defined in D and $\mathcal{D}[\![D]\!]_{\mathcal{S}} \dot{\rightarrow} \mathcal{S}$, then *the program* $\mathcal{F}[\![D]\!]$ satisfies \mathcal{S} .

In any other situation, a positive result of the abstract diagnosis corresponds to proving that the program *strongly* satisfies the temporal property. Also if we provide a specification that is in contradiction with the actual behavior of the process, our technique answers as expected in a verification context.

5 An automatic decision procedure for csLTL

In order to make our abstract diagnosis approach effective, we need to define an automatic decision procedure to check the validity of the csLTL formulas that show up when checking a property. In particular, we need to handle csLTL formulas of the form $\psi \dot{\rightarrow} \phi$, where ψ corresponds to the computed approximated behavior of the program, and ϕ is the abstract intended behavior of the process.

We impose a restriction on the specification ϕ : we do not allow the use of existential quantifications. Actually, this restriction is quite natural in our context since, in general, we are interested in proving properties related to the

visible behavior of the program, not to the local variables. In contrast, negation can be applied to any formula ϕ (not only to constraints).

In this section, we extend the tableau construction for Propositional LTL (PLTL) of [11,13] in order to deal with csLTL formulas. We need to adapt the method to our context due to three issues:

1. The structures on which the logic is interpreted are different. In our case, traces (sequences of states) are monotonic, meaning that the information in each state always increases.
2. The logic itself is a bit different from PLTL since propositions are replaced by constraints in \mathbf{C} .
3. We have to handle existential quantification over variables of the underlying constraint system. This does not mean that we are dealing with a first-order logic as will become clear later

In the following, we first present the basic rules that are used during the construction of the tree associated to the tableau. Then we present the algorithm that implements the process of construction of the tree.

5.1 Basic rules for a csLTL tableau

Classic tableaux algorithms are based on the systematic construction of a graph which is used to check the satisfiability of the formula. In [11,13], the authors present an algorithm that does not need to use auxiliary structures such as graphs to decide about the satisfaction of the formula, and this makes this approach more suitable for automatization.

A tableau procedure is defined by means of rules that build a tree whose nodes are labeled with sets of formulas. If all branches of the tree are *closed*, then the formula has no models. Otherwise, we can obtain a model that satisfies the formula from the open branches. Let us introduce the basic rules for the csLTL case. As usual, we present just the minimal set of rules.

A tableau rule is applied to a node n labeled with the set of formulas $L(n)$. Each rule application requires a previous selection of a formula ϕ from $L(n)$. We call context to the set of formulas $L(n) \setminus \{\phi\}$ and we denote it with Γ . Conjunctions are α -formulas and disjunctions β -formulas. Fig. 2 presents the rules for α - and β -formulas.

	α	$A(\alpha)$	β	$B_1(\beta)$	$B_2(\beta)$	
			R3	$\neg(\phi_1 \wedge \phi_2)$	$\{\neg\phi_1\}$	$\{\neg\phi_2\}$
R1	$\neg\neg\phi$	$\{\phi\}$	R4	$\neg(\phi_1 \mathcal{U} \phi_2)$	$\{\neg\phi_1, \neg\phi_2\}$	$\{\phi_1, \neg\phi_2, \neg\bigcirc(\phi_1 \mathcal{U} \phi_2)\}$
R2	$\phi_1 \wedge \phi_2$	$\{\phi_1, \phi_2\}$	R5	$\phi_1 \mathcal{U} \phi_2$	$\{\phi_2\}$	$\{\phi_1, \neg\phi_2, \bigcirc(\phi_1 \mathcal{U} \phi_2)\}$
			R6	$\phi_1 \mathcal{U} \phi_2$	$\{\phi_2\}$	$\{\phi_1, \neg\phi_2, \bigcirc((\Gamma^* \wedge \phi_1) \mathcal{U} \phi_2)\}$

Fig. 2. α - and β -formulas rules

Tables in Fig. 2 are interpreted as follows. Each row in a table represents a rule. Each time that an α -rule is applied to a node of the tree, a formula of the node matching the pattern in column α is replaced in a child node by the corresponding $A(\alpha)$. For the β -rules, two children nodes are generated, one for each column $B_1(\beta)$ and $B_2(\beta)$.

Almost all the rules are standard. However, Rule R6 uses the so-called context Γ^* , which is defined as $\Gamma^* := \dot{\forall}_{\gamma \in \Gamma} \dot{\neg} \gamma$. The use of contexts is the mechanism to detect the loops where no formula changes, thus allowing to mark branches containing eventually formulas as *open*. This kind of rules were first used in [12].

Note that there is no rule defined for the \bigcirc operator. In fact, the $\text{next}(\Phi)$ function transforms a set of elementary formulas Φ into another: $\text{next}(\Phi) := \{\phi \mid \bigcirc \phi \in \Phi\} \cup \{\dot{\neg} \phi \mid \dot{\neg} \bigcirc \phi \in \Phi\} \cup \{c \mid c \in \Phi, c \in \mathbf{C}\}$. This operator is different from the corresponding one of PLTL in that, in addition to keeping the internal formula of the next formulas, it also *passes* the constraints that are entailed at the current time instant to the following one. This makes sense for *tccp* computations since, as already mentioned, the store in a computation is monotonic, thus no information can be removed and it happens that, always, c implies $\bigcirc c$.

The next operator is a key notion in the kind of tableaux defined in [11,13]. This operator allows one to identify *stages* in a tableau which represent time instants in the model.

We show that α - and β -formulas rules and the next operator preserve the satisfiability of a set of formulas.

Lemma 1. *Given a set of formulas Φ , an α -formula α and a β -formula β :*

1. $\Phi \cup \{\alpha\}$ is satisfiable $\Leftrightarrow \Phi \cup A(\alpha)$ is satisfiable;
2. $\Phi \cup \{\beta\}$ is satisfiable $\Leftrightarrow \Phi \cup B_1(\beta)$ or $\Phi \cup B_2(\beta)$ is satisfiable;
3. if Φ is a set of elementary formulas, Φ is satisfiable $\Leftrightarrow \text{next}(\Phi)$ is satisfiable;

A second main difference w.r.t. the PLTL case regards the existential quantification. The *csLTL* existential quantification is introduced to model information about local variables, thus, the formula $\dot{\exists}_x \phi$ can be seen as the formula ϕ where the information about x is local.

We define a specific rule for the $\dot{\exists}$ case: when the selected formula of a given node is of the form $\dot{\exists}_x \phi$, it is created a node, child of n , whose labeling is that of n except that the formula $\dot{\exists}_x \phi$ is replaced by ϕ . Correctness of this rule derives from the following lemma, which shows that $\dot{\exists}_x \phi$ and ϕ are equi-satisfiable.

Lemma 2. *Let $\phi \in \text{csLTL}$, $\dot{\exists}_x \phi$ is satisfiable $\Leftrightarrow \phi$ satisfiable.*

5.2 Semantic *csLTL* tableaux

In this section, we present the notion of tableau for our *csLTL* formulas following the ideas of [11,13]. Since we borrow some definitions and notions from that work, in this section we skip some formal definitions.

A tableau \mathcal{T}_Φ for a set of formulas Φ is a tree-like structure where each node n is labeled with a set of *csLTL* formulas $L(n)$. The root is labeled with the

set of formulas Φ whose satisfiability/unsatisfiability is needed to check; Then, children of nodes are the result of applying the basic rules of Subsection 5.1. The algorithm in which these nodes are built is given in the following subsection. Nodes with no children are called *leaf* nodes.

Definition 3 (csLTL tableau). A csLTL tableau for a finite set of formulas Φ is a tuple $\mathcal{T}_\Phi = (Nodes, n_\Phi, L, B, R)$ such that:

1. *Nodes* is a finite non-empty set of nodes;
2. $n_\Phi \in Nodes$ is the initial node;
3. $L : Nodes \rightarrow \wp(csLTL)$ is the labeling function that associates to each node the formulas which are true in that node; the initial node is labeled with Φ ;
4. B is the set of branches such that exactly one of the following points holds for every $b = n_0, \dots, n_i, n_{i+1}, \dots, n_k \in B$ and every $0 \leq i < k$:
 - (a) $L(n_{i+1}) = \{A(\alpha)\} \cup L(n_i) \setminus \{\alpha\}$ for an α -formula $\alpha \in L(n_i)$;
 - (b) $L(n_{i+1}) = \{B_1(\beta)\} \cup L(n_i) \setminus \{\beta\}$ and there exists $b' = n_0, \dots, n_i, n'_{i+1}, \dots, n'_k$ such that $L(n'_{i+1}) = \{B_2(\beta)\} \cup L(n_i) \setminus \{\beta\}$ for a β -formula $\beta \in L(n_i)$;
 - (c) $L(n_{i+1}) = \{\phi'\} \cup L(n_i) \setminus \{\exists_x \phi'\}$ for $\exists_x \phi' \in L(n_i)$;
 - (d) $L(n_{i+1}) = next(L(n_i))$ if $L(n_i)$ contains only elementary formulas.

A branch $b \in B$ is maximal if it is not a proper prefix of another branch in B .

Definition 4. A node in the tableau is inconsistent if it contains

- a couple of formulas $\phi, \dot{\neg} \phi$, or
- the formula false, or
- a couple of constraint formulas $c, \dot{\neg} c'$ such that $c \vdash c'$.

The last condition for inconsistency of a node is particular to the *ccp* context. Since we are dealing with constraints that model partial information, it is possible to have an *implicit* inconsistency, in the sense that we need the entailment relation to detect it.

An inconsistent node does not accept any rule application. When a branch contains an inconsistent node, it is said to be closed, otherwise it is open.

By Lemma 1 and by Definition 3, it can be noticed that every closed branch contains only unsatisfiable sets of formulas. Open branches are not necessarily satisfiable since they could be prefixes of a closed one.

Similarly to the PLTL case, it exists only a finite number of different labels in a tableau. Thus, an infinite branch $b = n_0, n_1, \dots, n_k \dots$ contains a cycle. These branches are called *cyclic branches* and can be finitely represented as $path(b) = n_0, n_1, \dots, n_j, (n_{j+1}, \dots, n_k)^\omega$ when $L(n_k) = L(n_j)$ for $0 \leq j < k$.

Every branch of a tableau is divided into stages. A *stage* is a sequence of consecutive nodes between two consecutive applications of the next operator. We abuse of notation and we say that the labeling of a stage s is the labeling of each node in that stage ($L(s) = \bigcup_{n \in s} L(n)$). Moreover, a stage s is *saturated* if no α -, β - or hiding rule can be applied to any of its nodes.

We borrow from [11,13] the characterization of *fulfilled* eventually formula in a path of the tableau, namely when it is satisfied. We say that, when an

eventually formula $\phi_1 \mathcal{U} \phi_2$ belongs to the labeling of a stage s in a path, it is fulfilled if there exists a subsequent stage s' such that $\phi_2 \in L(n')$. A sequence of stages S is fulfilling if all the eventually formulas in its labeling are fulfilled in S and a branch b is fulfilling if the sequences of stages in its paths are fulfilling.

Finally, an open branch is expanded if it is fulfilling and all its stages are saturated. These notions are needed to formalize the tableau construction since only branches that are non-expanded and open are selected to be further developed.

Definition 5 (expanded csLTL tableau [11,13]). *A tableau is called expanded if every branch is expanded or closed. An expanded tableau is closed if every branch ends in an inconsistent node, otherwise it is open.*

5.3 A systematic csLTL tableaux construction

Definition 6 presents the algorithm to automatically build an expanded csLTL tableau (called systematic tableau [11,13]) for a given set of formulas Φ .

The construction consists in selecting at each step a non-expanded branch to be enlarged by using α or β rules or \exists elimination. When none of these can be applied, the next operator is used to pass to the next stage.

When dealing with eventualities, to determine which rule R5 or R6 has to be applied in a node, it is necessary to *distinguish* the eventuality that is being unfolded in the path. In this way, the rule R6 is applied only to *distinguished* eventualities when selected; when the selected eventuality is not the distinguished one, then rule R5 is used. If a node does not contain any distinguished eventuality, then the algorithm distinguishes one of them and rule R6 is chosen to be applied to it. Each node of the tableau has at most one distinguished eventuality.

The algorithm marks a node as *closed* when it is inconsistent and as *open* when it contains just constraint formulas or when it is the last node of an expanded branch (all the eventualities in the path are fulfilled).

Definition 6. *Given as input a finite set of formulas Φ , the algorithm repeatedly selects an unmarked leaf node l labelled with the set of formulas $L(l)$ and applies, in order, one of the points shown below.*

1. If l is an inconsistent node, then mark it as closed (\times).
2. If $L(l)$ is a set of constraint formulas, mark l as open (\odot).
3. If $L(l) = L(l')$ for l' ancestor of l , take the oldest ancestor l'' of l that is labeled with $L(l)$ and check if each eventuality in the path between l'' and l is fulfilled in such path. If they are all fulfilled, then mark l as open (\odot).
4. Otherwise, choose $\phi \in L(l)$ such that ϕ is not a next formula. Then,
 - if ϕ is an α -formula (let $\phi = \alpha$), create a new node l' as a child of l and label it as $L(l') = (L(l) \setminus \{\alpha\}) \cup A(\alpha)$ by using the α -rules in Fig. 2,
 - if ϕ is a β -formula (let $\phi = \beta$), create two new nodes l' and l'' as children of l and label them as $L(l') = (L(l) \setminus \{\beta\}) \cup B_1(\beta)$ and $L(l'') = (L(l) \setminus \{\beta\}) \cup B_2(\beta)$ by using the corresponding rule in Fig. 2. Moreover, if β is an eventuality, we have three possible cases:

- if β is the distinguished eventuality in $L(l)$, then apply Rule R6 to β and distinguish the formula inside the next formula in $B_2(\beta)$;
 - if β is not distinguished, but there is another distinguished formula, then apply Rule R5 to β and maintain the existing distinguished formula in $B_1(\beta)$ and $B_2(\beta)$;
 - otherwise, distinguish β and apply Rule R6 to β and distinguish the formula inside the next formula in $B_2(\beta)$;
- if ϕ is an \exists -formula ($\phi = \exists_x \phi'$), then create a new node l' as a child of l and label it as $L(l') = (L(l) \setminus \{\phi\}) \cup \{\phi'\}$
5. If $L(l)$ is a set of elementary formulas, apply the next operator: create a new node l' as child of l and label it as $L(l') = \text{next}(L(l))$.

The construction terminates when every branch is marked.

By construction, each stage in the systematic tableau \mathcal{T}_Φ for Φ is saturated. In order to ensure termination of the algorithm, it is necessary to use a *fair* strategy to distinguish eventualities, in the sense that every eventuality in an open branch must be distinguished at some point. This assumption and the fact that, given a finite set of initial formulas, there exist only a finite set of possible labels in a systematic tableau, imply termination.

It is worth noticing that, by the application of the rules in Fig. 2, when both ϕ and $\neg\phi$ belong to the labeling of a stage in a branch b , then any branch prefixed by b is closed. Moreover, by construction, non-fulfilled undistinguished eventualities in a branch are kept until they are fulfilled or they become distinguished.

One key result of the tableau in [13] and that we borrow is that if a distinguished eventuality is not fulfilled in an expanded branch b , then we can mark the branch as closed. This is because if we apply Rule R6, then we get a contradiction with the context of the eventuality.

As a conclusion, we have that every distinguished eventuality in a cyclic branch b of \mathcal{T}_Φ is fulfilled just since if it were unfulfilled, then b would be closed (thus not cyclic). Also, by construction and the above properties, b is open if and only if (1) the last node of b contains only constraint formulas, or (2) b is cyclic and all its eventualities are fulfilled in b .

Lemma 3. *By using a fair strategy and given as input a finite set $\Phi \subseteq \text{csLTL}$, the algorithm of Definition 6, terminates and builds an expanded tableau for \mathcal{T}_Φ .*

5.4 Soundness and completeness

Let us now show that the proposed algorithm is sound and complete for proving the satisfiability/unsatisfiability of csLTL formulas.

Theorem 2 (soundness). *If there exists a closed systematic tableau for $\Phi \subseteq \text{csLTL}$, then Φ is unsatisfiable.*

In order to prove completeness, we need to define an auxiliary function *stores* that, given a sequence of stages, builds a suitable conditional trace which joins

all the accumulated information in a stage at each time instant. We abuse of notation and write ϵ the empty sequence of stages. Recall that \otimes is the join operation of the constraint system and $\otimes \emptyset = \text{true}$.

$$\begin{aligned} \text{stores}(\epsilon) &= \epsilon \\ \text{stores}(s \cdot S) &= (C, \emptyset) \mapsto C \cdot \text{stores}(S) \quad \text{where } C = \otimes \{c \mid c \in \mathbf{C}, c \in L(n), n \in s\} \end{aligned}$$

By definition of `next`, which in our case propagates the constraints from one stage to the following, the conditional trace r resulting of applying `stores` to a sequence of stages S is monotone. Furthermore, since all the negative conditions are empty, r is also consistent.

We show that, given a systematic tableau \mathcal{T}_Φ built for Φ , we can compute a model for Φ from every open branch b in \mathcal{T}_Φ .

Lemma 4. *Let b be an open expanded branch in the systematic tableau \mathcal{T}_Φ for $\Phi \subseteq \text{csLTL}$. Given the sequence of stages S in $\text{path}(b)$, then $\text{stores}(S) \models \Phi$.*

Theorem 3 (refutational completeness). *If $\Phi \subseteq \text{csLTL}$ is unsatisfiable, then there exists a closed tableau for Φ .*

Theorem 4 (completeness). *If $\Phi \subseteq \text{csLTL}$ is satisfiable, then there exists a finite open tableau for Φ .*

5.5 Application of the Tableau

We are interested in checking the validity of a formula of the form $\psi \dot{\rightarrow} \phi$. Our strategy is to build the tableau for its negation $\mathcal{T}_{\neg(\psi \dot{\rightarrow} \phi)}$ so that, if $\mathcal{T}_{\neg(\psi \dot{\rightarrow} \phi)}$ is closed, meaning that $\neg(\psi \dot{\rightarrow} \phi)$ is unsatisfiable, then our implication is valid.

Let us show two examples of construction of the systematic tableaux for two formulas of this kind.

Example 4. Consider the formula of our guiding example $\exists_x \phi \dot{\rightarrow} \diamond(y = 1)$, where $\phi = (y = 1 \wedge \bigcirc x = 5 \wedge \bigcirc(\diamond y = 1)) \dot{\vee} (\neg y = 1 \wedge \bigcirc y = 1)$. The tableau in Fig. 3, with $L(\text{root}) = \exists_x \phi \wedge \square \neg(y = 1)$ shows its validity. Arrows labeled with α and β correspond to the application of α and β rules, respectively; arrows labeled with X represent the application of the `next` operator. Finally, arrows labeled with \exists correspond to the elimination of the existential quantification.

In the example, the first step uses the rule for the conjunction. Then, the second step involves the elimination of the existential quantification for $\exists_x \phi$. Since the formula $\square(\neg y = 1)$, which represents the context, does not contain information about x , $\exists_x \phi$ can be replaced with ϕ . The formula ϕ is then selected for a β step (disjunction). The branch on the left is closed after two steps since $y = 1$ and $\neg y = 1$ both belong to the node labeling.

The branch on the right, first flattens the conjunction and then applies the `next` rule. Note that the negation of a constraint is not kept in the following time instant. We recall that negation means “not entailment” (in contrast to meaning that *the contrary is true*), thus, in the future, the constraint could become true.

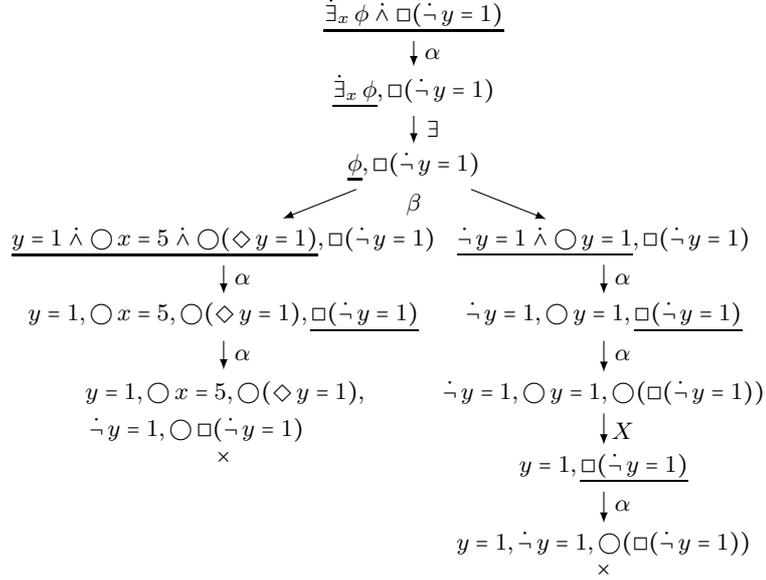


Fig. 3. Tableau for $\exists_x \phi \rightarrow \diamond y = 1$ of Example 4.

Since both branches are closed, we know that the formula $\exists_x \phi \wedge \Box(\neg(y = 1))$ is not satisfiable, thus its negation $\exists_x \phi \rightarrow \diamond(y = 1)$ is valid.

In the context of abstract diagnosis, this proves that the program is abstractly correct w.r.t. the LTL specification.

Example 5. Now, suppose that we want to check, for the program introduced in Example 1, that the constraint $y = 1$ is always entailed by the store. The corresponding specification is $\mathcal{S}'(p(y)) = \Box(y = 1)$.

The csLTL-semantics \mathcal{D} for $p(y)$ using the given specification as interpretation is $\exists_x((y = 1 \wedge \bigcirc x = 5 \wedge \bigcirc(\Box y = 1)) \dot{\vee} (\neg y = 1 \wedge \bigcirc y = 1))$. Let us abbreviate the body of the existential quantification as ϕ' . To check whether the process $p(y)$ is correct w.r.t. the property, we have to show that $\exists_x \phi' \rightarrow \Box(y = 1)$ is valid.

Fig. 4 shows part of the (finite) tableau that proves the satisfiability of the formula $\exists_x \phi' \wedge \diamond(\neg y = 1)$. This means that its negation, $\exists_x \phi' \rightarrow \Box(y = 1)$, is not valid. In the context of abstract diagnosis, although the formula is actually not satisfied by the program, because of the loss of precision due to the approximation, this is only a warning about the possible incorrectness of the program w.r.t. the LTL specification.

6 How to handle streams

In *tccp* streams are used to model imperative-style variables. In this context, when specifying an intended behavior, we are interested in the *current* value in a stream S (i.e., the last instantiated value in the stream), denoted as $S \doteq \text{value}$.

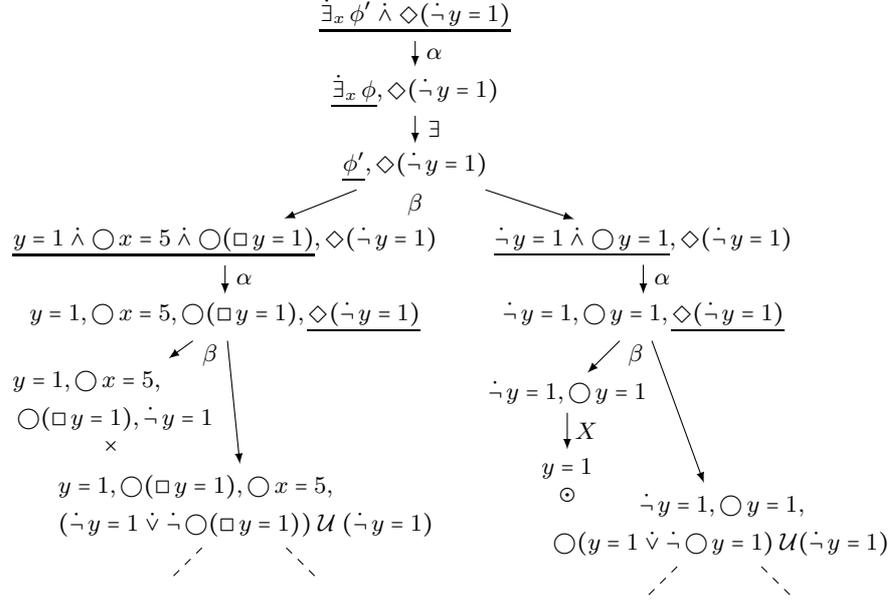


Fig. 4. Tableau for $\exists x \phi' \rightarrow \bigcirc y = 1$ of Example 5

In order to deal in our tableau with this special kind of constraint, we define a stream simplification function σ that transforms a formula ϕ into a new formula ϕ' that contains information only about the last instantiated value of the streams in ϕ . We show that the transformation preserves satisfiability.

The function σ uses two auxiliary functions. $dep(\phi)$ generates a set of dependencies on the form (S, S') , where S is a stream and S' its tail in the formula ϕ , while $name(S, D)$ returns, given a set of dependencies D , the first name given to the stream S . For instance, for $\phi = ((S = [c \mid S']) \wedge (S' = [c' \mid S''])) \mathcal{U} (T = [d \mid T'])$, $dep(\phi) = \{(S, S'), (S', S''), (T, T')\}$ and $name(S'', dep(\phi)) = S$. Formally,

$$dep(\phi) := \begin{cases} \emptyset & \text{if } \phi = \text{true}, \phi = \text{false} \text{ or } \phi = c \\ (S, S') & \text{if } \phi = (S = [c \mid S']) \\ dep(\phi_1) & \text{if } \phi = \exists x \phi_1 \text{ or } \phi = \bigcirc \phi_1 \\ dep(\phi_1) \cup dep(\phi_2) & \text{if } \phi = \phi_1 \wedge \phi_2 \text{ or } \phi = \phi_1 \mathcal{U} \phi_2 \end{cases}$$

$$name(S, D) := \begin{cases} S & \text{if } D = \emptyset \\ name(S', D') & \text{if } D = \{(S', S)\} \cup D' \\ name(S, D') & \text{if } D = \{(T', T)\} \cup D' \text{ and } T \neq S \end{cases}$$

Definition 7. Let $\phi, \phi_1, \phi_2 \in \text{csLTL}$, $c \in \mathbf{C}$ and S, S', R, R' streams of \mathbf{C} . σ is defined inductively as follows.

$$\sigma(\phi) := \begin{cases} \phi & \text{if } \phi = \text{true}, \phi = \text{false} \text{ or } \phi = c \\ S \dot{=} c & \text{if } \phi = (S' = [c \mid S'']) \text{ and } \text{name}(S', \text{dep}(\phi)) = S \\ \dot{\neg} \sigma(\phi_1) & \text{if } \phi = \dot{\neg} \phi_1 \\ \sigma(\phi_1) \wedge \sigma(\phi_2) & \text{if } \phi = \phi_1 \wedge \phi_2 \\ \dot{\exists}_x \sigma(\phi_1) & \text{if } \phi = \dot{\exists}_x \phi_1 \\ \bigcirc \sigma(\phi_1) & \text{if } \phi = \bigcirc \phi_1 \\ \sigma(\phi_1) \mathcal{U} \sigma(\phi_2) & \text{if } \phi = \phi_1 \mathcal{U} \phi_2 \end{cases}$$

Example 6. Consider $\phi = (C = [\text{near} \mid C'] \wedge \bigcirc(C' = [\text{out} \mid C''])) \wedge G = [\text{down} \mid G']$, then $\sigma(\phi) = (C \dot{=} \text{near} \wedge \bigcirc(C \dot{=} \text{out})) \wedge G \dot{=} \text{down}$.

The transformation σ preserves satisfiability, thus it can be used to preprocess the initial formula before applying the tableau method.

Lemma 5. Let $\phi \in \text{csLTL}$, ϕ is satisfiable $\iff \sigma(\phi)$ is satisfiable.

If we need to use this transformation, then the next operator must be slightly modified: $\text{next}(\Phi) := \{\phi \mid \bigcirc \phi \in \Phi\} \cup \{\dot{\neg} \phi \mid \dot{\neg} \bigcirc \phi \in \Phi\} \cup \{c \mid c \in \Phi, c \in \mathbf{C}\} \cup \{S \dot{=} c_1 \mid S \dot{=} c_1 \in \Phi \text{ and } \nexists c_2 \text{ such that } \bigcirc(S \dot{=} c_2) \in \Phi \text{ and } c_1 \neq c_2\}$. Intuitively, the constraints on the form $S \dot{=} c_1$ are propagated only if in the next time instant the tail of S has not be instantiated with a different constraint c_2 . This definition of next preserves satisfiability of Lemma 1.

7 Conclusions and Future Work

In this paper, we have introduced a decision procedure for csLTL formulas. csLTL is a linear temporal logic that replaces propositional formulas by constraint formulas, thus in order to determine the validity of a formula with no temporal constructs, it uses the entailment relation of the underlying constraint system.

This decision procedure is the last step of a method to validate LTL formulas for *tccp* programs. It is an adaptation of the tableau defined in [11,13]. The main differences of our algorithm w.r.t. the propositional case are due to the constraint nature of the behavior of *tccp*.

A Constraint Linear Temporal Logic is defined in [18] for the verification of a different timed concurrent language, called *ntcc*, which shares with *tccp* the concurrent constraint nature and the non-monotonic behavior. A fragment of the proposed logic, the restricted negation fragment where negation is only allowed for state formulas, is shown to be decidable. However, no efficient decision procedure is given. Moreover, the verification results are given for the locally-independent fragment of *ntcc*, which avoids the non-monotonicity of the original language. In contrast, our abstract diagnosis technique checks temporal properties for the full *tccp* language.

As future work, we plan to implement this algorithm and integrate it with the verification method. We also plan to explore other instances of the method based on logics for which decision procedures or (semi)automatic tools exists.

References

1. M. Alpuente, M. Comini, S. Escobar, M. Falaschi, and S. Lucas. Abstract Diagnosis of Functional Programs. *LOPSTR 2002, Revised Selected Papers*, volume 2664 of *LNCS*, pages 1–16, 2003. Springer-Verlag.
2. G. Bacci and M. Comini. Abstract Diagnosis of First Order Functional Logic Programs. *LOPSTR 2010*, volume 6564 of *LNCS*, pages 215–233, 2011. Springer-Verlag.
3. E. M. Clarke and E. A. Emerson. Design and synthesis of synchronization skeletons using branching-time temporal logic. *Logic of Programs*, volume 131 of *LNCS*, pages 52–71, 1981. Springer-Verlag.
4. M. Comini, G. Levi, M. C. Meo, and G. Vitiello. Abstract Diagnosis. *Journal of Logic Programming*, 39(1-3):43–93, 1999.
5. M. Comini, L. Titolo, and A. Villanueva. Abstract Diagnosis for Timed Concurrent Constraint programs. *TPLP*, 11(4-5):487–502, 2011.
6. M. Comini, L. Titolo, and A. Villanueva. A Condensed Goal-Independent Bottom-Up Fixpoint Modeling the Behavior of *tccp*. Technical report, DSIC, UPV. Available at <http://riunet.upv.es/handle/10251/8351>, 2013.
7. F. S. de Boer, M. Gabbrielli, and M. C. Meo. A Timed Concurrent Constraint Language. *Information and Computation*, 161(1):45–83, 2000.
8. F. S. de Boer, M. Gabbrielli, and M. C. Meo. A Temporal Logic for Reasoning about Timed Concurrent Constraint Programs. *TIME '01*, page 227, Washington, DC, USA, 2001. IEEE Computer Society.
9. F. S. de Boer, M. Gabbrielli, and M. C. Meo. Proving correctness of Timed Concurrent Constraint Programs. *CoRR*, cs.LO/0208042, 2002.
10. M. Falaschi, C. Olarte, C. Palamidessi, and F. D. Valencia. Declarative Diagnosis of Temporal Concurrent Constraint Programs. *ICLP 2007*, volume 4670 of *LNCS*, pages 271–285. Springer-Verlag, 2007.
11. J. Gaintzarain, M. Hermo, P. Lucio, and M. Navarro. Systematic semantic tableaux for PLTL. *Electronic Notes in Theoretical Computer Science*, 206:59–73, 2008.
12. J. Gaintzarain, M. Hermo, P. Lucio, M. Navarro, and F. Orejas. A cut-free and invariant-free sequent calculus for PLTL. *CSL*, volume 4646 of *LNCS*, pages 481–495. Springer, 2007.
13. J. Gaintzarain, M. Hermo, P. Lucio, M. Navarro, and F. Orejas. Dual systems of tableaux and sequents for PLTL. *JLAP*, 78(8):701–722, 2009.
14. Z. Manna and A. Pnueli. *The temporal logic of reactive and concurrent systems - specification*. Springer, 1992.
15. C. Palamidessi and F. D. Valencia. A Temporal Concurrent Constraint Programming Calculus. *CP'01*, volume 2239 of *LNCS*, pages 302–316. Springer, 2001.
16. J. P. Queille and J. Sifakis. Specification and verification of concurrent systems in CESAR. *Symposium on Programming*, volume 137 of *LNCS*, pages 337–351. Springer, 1982.
17. V. A. Saraswat. *Concurrent Constraint Programming*. The MIT Press, Cambridge, Mass., 1993.
18. F. D. Valencia. Decidability of infinite-state timed ccp processes and first-order LTL. *Theoretical Computer Science*, 330(3):577–607, 2005.