

*Abstract diagnosis for timed concurrent constraint programs**

MARCO COMINI and LAURA TITOLO

*Dipartimento di Matematica e Informatica University of Udine
Via delle Scienze, 206 33100 Udine, Italy
(e-mail: {marco.comini,laura.titolo}@uniud.it)*

ALICIA VILLANUEVA

*Departamento de Sistemas Informáticos y Computación Universitat Politècnica de València
Camino de Vera s/n 46022 Valencia, Spain
(e-mail: villanue@dsic.upv.es)*

Abstract

The timed concurrent constraint language (tccp in short) is a concurrent logic language based on the simple but powerful concurrent constraint paradigm of Saraswat. In this paradigm, the notion of store-as-value is replaced by the notion of store-as-constraint, which introduces some differences w.r.t. other approaches to concurrency. In this paper, we provide a general framework for the debugging of tccp programs. To this end, we first present a new compact, bottom-up semantics for the language that is well suited for debugging and verification purposes in the context of reactive systems. We also provide an abstract semantics that allows us to effectively implement debugging algorithms based on abstract interpretation. Given a tccp program and a behavior specification, our debugging approach automatically detects whether the program satisfies the specification. This differs from other semi-automatic approaches to debugging and avoids the need to provide symptoms in advance. We show the efficacy of our approach by introducing two illustrative examples. We choose a specific abstract domain and show how we can detect that a program is erroneous.

KEYWORDS: concurrent constraint paradigm; denotational semantics; abstract diagnosis; abstract interpretation

1 Introduction

Finding program bugs is a long-standing problem in software construction. In the concurrent paradigms, the problem is even worse and the traditional tracing techniques are almost useless. There has been a lot of work on algorithmic debugging (Shapiro 1982) for declarative languages, which could be a valid proposal for concurrent paradigms, but little effort has been done for the particular case of the

* This work has been partially supported by the EU (FEDER), the Spanish MICINN under grant TIN2010-21062-C02-02 and by the Universitat Politècnica de València under grant PAID-00-10.

concurrent constraint paradigm (ccp in short; Saraswat 1993). The ccp paradigm is different from other programming paradigms mainly due to the notion of store-as-constraint that replaces the classical store-as-valuation model. In this way, the languages from this paradigm can easily deal with partial information: an underlying constraint system handles constraints on system variables. Within this family de Boer *et al.* (2000) introduced the timed concurrent constraint language (tccp in short) by adding to the original ccp model the notion of time and the ability to capture the absence of information. With these features, it is possible to specify behaviors typical of reactive systems such as timeouts or preemption actions, but they also make the language nonmonotonic.

In this paper, we develop an abstract diagnosis method for tccp using the ideas of the abstract diagnosis framework for logic programming (Comini *et al.* 1999). This framework, parametric w.r.t. an abstract program property, is based on the use of an abstract immediate consequence operator to identify bugs in logic programs. It can be considered as an extension of algorithmic debugging since there are instances of the framework that deliver the same results. The intuition of the approach is that, given an abstract specification of the expected behavior of the program, one automatically detects the errors in the program. The framework does not require the determination of symptoms in advance. To achieve an effective method, abstract interpretation is used to approximate the semantics, thus results may be less precise than those obtained by using the concrete semantics.

The approach of abstract diagnosis for logic programming has been applied to other paradigms (Alpuente *et al.* 2003; Falaschi *et al.* 2007; Bacci and Comini 2011). This research revealed that a key point for the efficacy of the resulting debugging methodology is the compactness of the concrete semantics. Thus, in this proposal, much effort has been devoted to the development of a compact concrete semantics for the tccp language to start with. The already existing denotational semantics are based on capturing the input–output behavior of the system. However, since we are in a concurrent (reactive) context, we want to analyze and debug infinite computations. Our semantics covers this need and is suitable to be used not only with debugging techniques but also with other verification approaches.

Our new (concrete) compact compositional semantics is correct and fully abstract w.r.t. the small-step behavior of tccp. It is based on the evaluation of agents over a denotation for a set of process declarations D , obtained as least fixpoint of a (continuous, monotone) immediate consequence operator $\mathcal{D}[[D]]$.

Thanks to the compactness of this semantics, we can formulate an efficacious debugging methodology based on abstract interpretation which proceeds by approximating the $\mathcal{D}[[D]]$ operator producing an “abstract immediate consequence operator” $\mathcal{D}^\alpha[[D]]$. We show that, given the abstract intended specification \mathcal{S}^α of the semantics of the declarations D , we can check the correctness of D by a single application of $\mathcal{D}^\alpha[[D]]$ and thus, by a static test, we can determine all the process declarations $d \in D$ which are wrong w.r.t. the considered abstract property.

To our knowledge, in the literature there is only another approach to the debugging problem of ccp languages (Falaschi *et al.* 2007), which is also based on the abstract diagnosis approach of (Comini *et al.* 1999). However, they consider a quite different

concurrent constraint language without nonmonotonic features, which we consider essential to model behaviors of reactive systems.

2 The timed concurrent constraint language

The tcp language is particularly suitable to specify both reactive and time critical systems. As the other languages of the ccp paradigm (Saraswat 1993), it is parametric w.r.t. a cylindric constraint system. The constraint system handles the data information of the program in terms of constraints. In tcp, 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 that store. In brief, a cylindric constraint system¹ $\mathbf{C} = \langle \mathcal{C}, \leq, \otimes, \oplus, tt, ff, \text{Var}, \exists \rangle$ is composed of a set of finite constraints \mathcal{C} ordered by \leq , where \oplus and \otimes are the *glb* and *lub*, respectively. tt is the smallest constraint whereas ff is the largest one. We often use the inverse order \vdash (called *entailment*) instead of \leq over constraints. Var is a denumerable set of variables and \exists existentially quantifies variables over constraints (the so called cylindric operator).

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

$$A ::= \text{skip} \mid \text{tell}(c) \mid \sum_{i=1}^n \text{ask}(c_i) \rightarrow A_i \mid \text{now } c \text{ then } A_1 \text{ else } A_2 \mid A_1 \parallel A_2 \mid \exists x A \mid p(\vec{x}),$$

where c and c_i are finite constraints in \mathcal{C} , $p \in \Pi$, $x \in \text{Var}$ and \vec{x} is a list of variables x_1, \dots, x_n with $1 \leq i \leq n$, $x_i \in \text{Var}$. A tcp program P is an object of the form $D.A_0$, where A_0 is an agent, called initial agent, and D is a set of process declarations of the form $p(\vec{x}) :- A$ (for some agent A).

The notion of time is introduced by defining a discrete and global clock: it is assumed that the *ask* and *tell* agents take one time-unit to be executed. For the operational semantics of the language, the reader can consult (de Boer *et al.* 2000). 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 nondeterministically 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 satisfied by the store, the agent suspends. The agent *now* c then A else B behaves in the current time instant like A (respectively B) if c is (respectively is not) satisfied by the store. The satisfaction 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 B is executed both when $\neg c$ is satisfied, but also when neither c nor $\neg c$ are satisfied. $A \parallel B$ models the parallel composition of A and B in terms of maximal parallelism (in contrast to the

¹ See (Saraswat 1993; de Boer *et al.* 2000) for more details on cylindric constraint systems.

interleaving approach of ccp), i.e., all the enabled agents of A and B 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 \exists operator of the constraint system. Finally, the agent $p(\tilde{x})$ takes from D a declaration of the form $p(\tilde{x}) :- A$ and executes A at the following time instant. For the sake of simplicity, we assume that the set D of declarations is closed w.r.t. parameter names.

3 Modeling the small-step operational behavior of tccp

In this section, we introduce a denotational semantics that models the small-step behavior of tccp. Due to space limitations, in this paper we show the concrete semantics and the most relevant aspects of the abstract one. The missing definitions, as well as the proofs of all the results, can be found in Comini *et al.* (2011).

Let us formalize the notion of behavior for a set D of process declarations. It collects all the small-step computations associated to D as the set of (all the prefixes of) the sequences of computational steps, for all possible initial agents and stores.

Definition 1 (Small-step behavior of declarations)

Let D be a set of declarations, $Agent$ the set of possible agents, and \rightarrow the transition relation given by the operational semantics in de Boer *et al.* (2000). The small-step behavior of D is defined as follows:

$$\mathcal{B}^{ss}[D] := \bigcup_{\forall c \in \mathcal{C}, \forall A \in Agent} \mathcal{B}[D.A]_c,$$

where $\mathcal{B}[D.A]_c := \{c \cdot c_1 \cdot \dots \cdot c_n \mid \langle A, c \rangle \rightarrow \langle A_1, c_1 \rangle \rightarrow \dots \rightarrow \langle A_n, c_n \rangle\} \cup \{\epsilon\}$. We denote by \approx_{ss} the equivalence relation between declarations induced by \mathcal{B}^{ss} , namely $D_1 \approx_{ss} D_2 \Leftrightarrow \mathcal{B}^{ss}[D_1] = \mathcal{B}^{ss}[D_2]$.

The pair $\langle A_i, c_i \rangle$ denotes a configuration where A_i is the agent to be executed, and c_i the store at that computation step. Thus, the small-step behavior is the set of sequences of stores that are computed by the operational semantics of the language.

There are many languages where a compact compositional semantics has been founded on collecting the possible traces for the weakest store, since all traces relative to any other initial store can be derived by instance of the formers. In tccp, this does not work since the language is not monotonic: if we have all traces for an agent A starting from an initial store c and we execute A with a more instantiated initial store d , then new traces, not instances of the formers, can appear.

Furthermore, note that, since we are interested in a bottom-up approach, we cannot work assuming that we know the initial store. However, when we define the semantics of a conditional or choice agent where some guard must be checked, we should consider different execution branches depending on the guard satisfiability. To deal with all these particular features, our idea is that of associating conditions to computation steps, and to collect all possible minimal hypothetical computations.

3.1 The semantic domain

In de Boer *et al.* (2000), reactive sequences are used as semantic domain for the top-down semantics. These sequences are composed of a pair of stores $\langle c, c' \rangle$ for each time instant meaning that, given the store c , the program produces in one time instant the store c' . The store is monotonic, thus c' always contains more (or equal) information than c .

As we have explained before, this information is not enough for a bottom-up approach.² Our idea is to enrich the reactive sequence notion so that we keep information about the essential conditions that the store must satisfy to make the program proceed. We define a *condition* η as a pair $\eta = (\eta^+, \eta^-)$ where $\eta^+ \in \mathcal{C}$ (respectively $\eta^- \in \wp(\mathcal{C})$) is called positive (respectively negative) component. A condition is said to be *inconsistent* when its positive component is ff or when it entails any constraint in the negative component. Given a store $c \in \mathcal{C}$, we say that c satisfies η (written $c \triangleright \eta$) when c entails η^+ , $\eta^+ \neq ff$ and c does not entail any constraint from η^- . An inconsistent condition is satisfied by no store, while the pair (tt, \emptyset) is satisfied by any store.

A conditional reactive sequence is a sequence of conditional tuples, which can be of two forms: (i) a triple $\eta \rightarrow \langle a, b \rangle$ that is used to represent a computational step, i.e., the global store a becomes b at the next time instant only if $a \triangleright \eta$, or (ii) a construct $stutt(C)$ that models the suspension of the computation due to an ask agent, i.e., it represents the fact that there is no guard in C (the guards of the choice agent) entailed by the current store. We need this construct to distinguish a suspended computation from an infinite loop that does not modify the store.

Our denotations are composed of conditional reactive sequences:

Definition 2 (Conditional reactive sequence)

A conditional reactive sequence is a sequence of conditional tuples of the form $t_1 \dots t_n \dots$, maybe ended with \square , such that: for each $t_i = \eta_i \rightarrow \langle a_i, b_i \rangle$, $b_i \vdash a_i$ for $i \geq 1$, and for each $t_j = \eta_j \rightarrow \langle a_j, b_j \rangle$ such that $j > i$, $a_j \vdash b_i$. The empty sequence is denoted with ϵ . $s_1 \cdot s_2$ denotes the concatenation of two conditional reactive sequences s_1, s_2 .

A set of conditional reactive sequences is *maximal* if none of its sequences is the prefix of another. By \mathbb{M} we denote the domain of sets of maximal conditional reactive sequences, whose order is induced from its prefix closure, namely $R_1 \sqsubseteq R_2 \Leftrightarrow prefix(R_1) \subseteq prefix(R_2)$. $(\mathbb{M}, \sqsubseteq, \sqcup, \sqcap, \perp, \top)$ is a complete lattice.

3.2 Semantics evaluation function for agents

To associate a denotation to a set of process declarations, we need first to define the semantics for agents. Let us now introduce the notion of interpretation.

² In a top-down approach, the (initial) current store is propagated, thus decisions regarding the satisfaction or not of a given condition can be taken immediately.

Definition 3 (Interpretations)

Let $\mathbf{MGC} := \{p(\vec{x}) \mid p \in \Pi, \vec{x} \text{ are distinct variables}\}$ be the set of *most general calls*. An *interpretation* is a function $\mathbf{MGC} \rightarrow \mathbf{M}$ modulo variance.³ Two functions $I, J : \mathbf{MGC} \rightarrow \mathbf{M}$ are *variants*, denoted by $I \cong J$, if for each $\pi \in \mathbf{MGC}$ there exists a variable renaming ρ such that $(I\pi)\rho = J(\pi\rho)$. The semantic domain \mathbf{II} is the set of all interpretations ordered by the point-wise extension of \sqsubseteq .

The application of an interpretation \mathcal{I} to a most general call π , denoted by $\mathcal{I}(\pi)$, is the application $I(\pi)$ of any representative I of \mathcal{I} which is defined exactly on π . For example, if $\mathcal{I} = (\lambda\varphi(x, y). \{(tt, \emptyset) \rightarrow \langle tt, x = y \rangle\}) / \cong$ then $\mathcal{I}(\varphi(u, v)) = \{(tt, \emptyset) \rightarrow \langle tt, u = v \rangle\}$.

The technical core of our semantics definition is the agent semantics evaluation function which, given an agent and an interpretation, builds the maximal conditional reactive sequences of the agent.

Definition 4 (Agents Semantics)

Given an agent A and an interpretation \mathcal{I} , the semantics $\mathcal{A}[[A]]_{\mathcal{I}}$ is defined by structural induction:

$$\begin{aligned} \mathcal{A}[[\text{skip}]]_{\mathcal{I}} &= \{\square\} \\ \mathcal{A}[[\text{tell}(c)]]_{\mathcal{I}} &= \{(tt, \emptyset) \rightarrow \langle tt, c \rangle \cdot \square\} \end{aligned} \quad (1)$$

$$\begin{aligned} \mathcal{A}[[\sum_{i=1}^n \text{ask}(c_i) \rightarrow A_i]]_{\mathcal{I}} &= \bigsqcup_{i=1}^n \{(c_i, \emptyset) \rightarrow \langle c_i, c_i \rangle \cdot (c_i \odot s) \mid s \in \mathcal{A}[[A_i]]_{\mathcal{I}}\} \sqcup \\ &\quad \bigsqcup \{stutt(\cup_{i=1}^n c_i) \cdot s \mid s \in \mathcal{A}[[\sum_{i=1}^n \text{ask}(c_i) \rightarrow A_i]]_{\mathcal{I}}, \forall i \in [1, n]. c_i \neq tt\} \end{aligned} \quad (2)$$

$$\begin{aligned} \mathcal{A}[[\text{now}(d) \text{ then } A \text{ else } B]]_{\mathcal{I}} &= \{(d, \emptyset) \rightarrow \langle d, d \rangle \cdot \square \mid \square \in \mathcal{A}[[A]]_{\mathcal{I}}\} \sqcup \\ &\quad \bigsqcup \{(c^+ \otimes d, c^-) \rightarrow \langle c \otimes d, c' \otimes d \rangle \cdot (d \odot s) \mid (c^+, c^-) \rightarrow \langle c, c' \rangle \cdot s \in \mathcal{A}[[A]]_{\mathcal{I}}, \\ &\quad \quad \quad c \otimes d \triangleright (c^+ \otimes d, c^-)\} \sqcup \end{aligned}$$

$$\bigsqcup \{(d, C) \rightarrow \langle d, d \rangle \cdot (d \odot s) \mid stutt(C) \cdot s \in \mathcal{A}[[A]]_{\mathcal{I}}, d \triangleright (d, C)\} \sqcup$$

$$\bigsqcup \{(tt, d) \rightarrow \langle tt, tt \rangle \cdot \square \mid \square \in \mathcal{A}[[B]]_{\mathcal{I}}\} \sqcup$$

$$\bigsqcup \{(c^+, c^- \cup \{d\}) \rightarrow \langle c, c' \rangle \cdot s \mid (c^+, c^-) \rightarrow \langle c, c' \rangle \cdot s \in \mathcal{A}[[B]]_{\mathcal{I}}, \\ c \triangleright (c^+, c^- \cup \{d\})\} \sqcup$$

$$\bigsqcup \{(tt, C \cup \{d\}) \rightarrow \langle tt, tt \rangle \cdot s \mid stutt(C) \cdot s \in \mathcal{A}[[B]]_{\mathcal{I}}\} \quad (3)$$

$$\mathcal{A}[[A \parallel B]]_{\mathcal{I}} = \bigsqcup \{s_A \parallel s_B \mid s_A \in \mathcal{A}[[A]]_{\mathcal{I}}, s_B \in \mathcal{A}[[B]]_{\mathcal{I}}\} \quad (4)$$

$$\mathcal{A}[[\exists x A]]_{\mathcal{I}} = \bigsqcup \{s \in \mathbf{M} \mid \exists s' \in \mathcal{A}[[A]]_{\mathcal{I}} \text{ such that } \exists_x s = \exists_x s', \\ s' \text{ is } x\text{-connected, } s \text{ is } x\text{-invariant}\} \quad (5)$$

$$\mathcal{A}[[p(z)]]_{\mathcal{I}} = \bigsqcup \{(tt, \emptyset) \rightarrow \langle tt, tt \rangle \cdot s \mid s \in \mathcal{I}(p(z))\}$$

Let us now illustrate the idea of the semantics. The tell agent works independently of the current store, thus in (1), the conditional reactive sequence starts with a conditional tuple composed by the condition (tt, \emptyset) , which is always satisfied, and a second part that says that the constraint c is added during the first computational step; afterwards, the computation terminates with \square .

³ That is, a family of elements of \mathbf{M} , indexed by \mathbf{MGC} , modulo variance.

The semantics for the nondeterministic choice (2), collects for each guard c_i a conditional sequence of the form $(c_i, \emptyset) \rightarrow \langle c_i, c_i \rangle \cdot (c_i \odot s)$. The condition states that c_i has to be satisfied by the current store, whereas the pair $\langle c_i, c_i \rangle$ represents the fact that the query to the store does not modify the store. The constraint c_i is propagated into the sequence s (the continuation of the computation which belongs to the semantics of A_i) by means of the propagation operator that (consistently) adds a given constraint to the stores appearing in a sequence:

$$h \odot s = \begin{cases} \eta \rightarrow \langle a \otimes h, b \otimes h \rangle \cdot (h \odot s') & \text{if } s = \eta \rightarrow \langle a, b \rangle \cdot s', \eta^+ \otimes h \neq ff, \\ & b \otimes h \neq ff, a \otimes h \triangleright \eta \\ \eta \rightarrow \langle a \otimes h, ff \rangle & \text{if } s = \eta \rightarrow \langle a, b \rangle \cdot s', \eta^+ \otimes h \neq ff, \\ & b \otimes h = ff, a \otimes h \triangleright \eta \\ stutt(\eta^-) \cdot (h \odot s') & \text{if } s = stutt(\eta^-) \cdot s' \\ s & \text{if } s = \epsilon \text{ or } s = \square \end{cases}$$


In addition, we have to model the case when the computation suspends, i.e., when no guard of the agent is satisfied by the current store. Sequences representing this situation are of the form $stutt(\bigcup_{i=1}^n \{c_i\}) \cdot s$ where s is, recursively, an element of the semantics of the choice agent. The only case when we do not include the stuttering sequence is when one of the guards c_i is tt . Note that, due to the partial nature of the constraint system, the fact that the disjunction of the guards is tt is not a sufficient condition to avoid suspension.

The definition of the conditional agent now is similar to the previous one. However, since it is instantaneous, we have six cases depending on the three possible heads of the sequences of the semantics of A (respectively B) and on the fact that the guard d is satisfied or not in the current time instant.

The semantics for the parallel composition of two agents (4), is defined in terms of an auxiliary commutative operator $\dot{\parallel}$ which combines the sequences of the two agents:

$$s_A \dot{\parallel} s_B = \begin{cases} (\eta \otimes_c \delta) \rightarrow \langle a \otimes c, b \otimes d \rangle \cdot (d \odot s'_A) \dot{\parallel} (b \odot s'_B) & \text{if } s_A = \eta \rightarrow \langle a, b \rangle \cdot s'_A, \\ & s_B = \delta \rightarrow \langle c, d \rangle \cdot s'_B, \\ & a \otimes c \triangleright (\eta \otimes_c \delta), b \otimes d \neq ff \\ (\eta \otimes_c \delta) \rightarrow \langle a \otimes c, ff \rangle & \text{if } s_A = \eta \rightarrow \langle a, b \rangle \cdot s'_A, \\ & s_B = \delta \rightarrow \langle c, d \rangle \cdot s'_B, \\ & a \otimes c \triangleright (\eta \otimes_c \delta), b \otimes d = ff \\ (\eta^+, \eta^- \cup \delta^-) \rightarrow \langle a, b \rangle \cdot s'_A \dot{\parallel} (b \odot s'_B) & \text{if } s_A = \eta \rightarrow \langle a, b \rangle \cdot s'_A, \\ & s_B = stutt(\delta^-) \cdot s'_B, \\ & a \triangleright (\eta^+, \eta^- \cup \delta^-) \\ stutt(\eta^- \cup \delta^-) \cdot s'_A \dot{\parallel} s'_B & \text{if } s'_A = stutt(\eta^-) \cdot s'_A, \\ & s'_B = stutt(\delta^-) \cdot s'_B \\ s_A & \text{if } s_B = \epsilon \text{ or } s_B = \square \end{cases}$$

For the hiding operator (5), we collect the sequences that satisfy the restrictions regarding the visibility of the hidden variables. In particular, a conditional reactive sequence $s = t_1 \dots t_n \dots$ is x -connected when (1) if $t_1 = \eta_1 \rightarrow \langle a_1, b_1 \rangle$ then $\exists_x a_1 = a_1$ and (2) for each $t_i = \eta_i \rightarrow \langle a_i, b_i \rangle$ and $t_{i+1} = \eta_{i+1} \rightarrow \langle a_{i+1}, b_{i+1} \rangle$, with $i > 1$,

Fig. 1. Tree representation of $\mathcal{F} \llbracket D \rrbracket$ in the example.

$\exists_x a_{i+1} \otimes b_i = a_{i+1}$. A conditional reactive sequence $s = t_1 \dots t_n \dots$ is *x-invariant* if for each computational step $t_i = \eta_i \rightarrow \langle a_i, b_i \rangle$, it holds that $b_i = \exists_x b_i \otimes a_i$.

Finally, the semantics of the process call $p(\vec{x})$ collects the sequences in the interpretation $\mathcal{J}(p(\vec{x}))$, delayed by one time unit, as stated in the operational semantics.

Let us show an illustrative example. Consider the tccp agent $A \equiv \text{ask}(y \geq 0) \rightarrow \text{tell}(z \leq 0)$. The semantics is composed of two sequences:

$$\begin{aligned} \mathcal{A} \llbracket A \rrbracket_{\mathcal{J}} = & \{ (y \geq 0, \emptyset) \rightarrow \langle y \geq 0, y \geq 0 \rangle \cdot (tt, \emptyset) \rightarrow \langle y \geq 0, y \geq 0 \otimes z \leq 0 \rangle \cdot \square \} \\ & \cup \{ \text{stutt}(y \geq 0) \cdot s \mid s \in \mathcal{A} \llbracket A \rrbracket_{\mathcal{J}} \}. \end{aligned}$$

3.3 Fixpoint denotations of declarations

Now we can define the semantics for a set of process declarations D as the fixpoint of the immediate consequences operator $\mathcal{D} \llbracket D \rrbracket_{\mathcal{J}} := \lambda p(x). \bigsqcup_{p(x) :- A \in D} \mathcal{A} \llbracket A \rrbracket_{\mathcal{J}}$, which is continuous. Thus, it has a least fixpoint and we can define the semantics of D as $\mathcal{F} \llbracket D \rrbracket = \text{lfp}(\mathcal{D} \llbracket D \rrbracket)$. As an example, in Figure 1 we represent the (infinite) set of traces of $\mathcal{F} \llbracket \{p(x) :- \exists y (\text{ask}(y > x) \rightarrow p(x + 1) + \text{ask}(y \leq x) \rightarrow \text{skip})\} \rrbracket$.⁴

In Comini et al. (2011) we have proven that $D_1 \approx_{ss} D_2$ if and only if $\mathcal{F} \llbracket D_1 \rrbracket = \mathcal{F} \llbracket D_2 \rrbracket$ (correctness and full abstraction of \mathcal{F} w.r.t. \approx_{ss}).

4 Abstract semantics for tccp: the abstraction scheme

In this section, starting from the fixpoint semantics in Section 3, we present an abstract semantics which approximates the observable behavior of the program. Program properties that are of interest are Galois Insertions between the concrete

⁴ For the sake of simplicity, we assume that we can use expressions of the form $x + 1$ directly in the arguments of a process call. We can simulate this behavior by writing $\text{tell}(x' = x + 1) \rightarrow p(x')$ (but introducing a delay of one time unit).

domain and the chosen abstract domain. We assume familiarity with basic results of abstract interpretation (Cousot and Cousot 1979).

We define an abstraction scheme where we develop the abstraction of computations, i.e., of maximal sets of conditional reactive sequences, by successive lifting. We start with a function that abstracts the information component of the program semantics, i.e., the store; then we build the abstraction of conditional tuples; then of conditional reactive sequences and, finally, of maximal sets.

We start from an upper-approximating function $\tau^+ : \mathcal{C} \rightarrow \hat{\mathcal{C}}$ into an abstract constraint system $\hat{\mathbf{C}} = \langle \hat{\mathcal{C}}, \hat{\preceq}, \hat{\otimes}, \hat{\oplus}, \hat{tt}, \hat{ff}, \text{Var}, \hat{\exists} \rangle$, where \hat{tt} and \hat{ff} are the smallest and the greatest abstract constraint, respectively. We often use the inverse relation $\hat{\vdash}$ of $\hat{\preceq}$. We have also a lower-approximating function $\tau^- : \wp(\mathcal{C}) \rightarrow \check{\mathcal{C}}$ into an abstract constraint system $\check{\mathbf{C}} = \langle \check{\mathcal{C}}, \check{\preceq}, \check{\otimes}, \check{\oplus}, \check{tt}, \check{ff}, \text{Var}, \check{\exists} \rangle$. This second function is needed to (correctly) deal with the negative part of conditions.

We have two “external” operations $\hat{\times} : \mathcal{C} \times \hat{\mathcal{C}} \rightarrow \hat{\mathcal{C}}$ and $\check{\times} : \mathcal{C} \times \check{\mathcal{C}} \rightarrow \check{\mathcal{C}}$ that update an abstract store with a concrete constraint (coming from the program). In addition, a “bridge” relation $\tilde{\vdash} \in \hat{\mathcal{C}} \times \check{\mathcal{C}}$ decides if an upper-abstract constraint is consistent with a lower-abstract constraint. Abstract and concrete constraint systems are related by these conditions:

$$\begin{array}{ll}
 c \hat{\times} \tau^+(a) = \tau^+(c \otimes a) & c \check{\times} \tau^-(C) = \tau^-(\{c\} \cup C) \\
 \tau^+(a \otimes b) = \tau^+(a) \hat{\otimes} \tau^+(b) & \tau^-(C \cup C') = \tau^-(C) \check{\oplus} \tau^-(C') \\
 a \vdash b \implies \tau^+(a) \hat{\vdash} \tau^+(b) & \tau^-(\{a\}) \check{\vdash} \tau^-(C) \implies \exists c \in C. a \vdash c \\
 \tau^+(\exists_x a) = \hat{\exists}_x \tau^+(a) & \tau^-(\{\exists_x c \mid c \in C\}) = \check{\exists}_x \tau^-(C) \\
 \forall c \in C. a \not\vdash c \implies \tau^+(a) \not\tilde{\vdash} \tau^-(C).
 \end{array}$$

An *abstract condition* is a pair of the form $(\hat{\eta}, \check{\eta}) \in \hat{\mathcal{C}} \times \check{\mathcal{C}}$. Similarly to the concrete case, given an abstract condition $\tilde{\eta} = (\hat{\eta}, \check{\eta})$ and an abstract store $\hat{a} \in \hat{\mathcal{C}}$, we say that \hat{a} satisfies $\tilde{\eta}$ (written $\hat{a} \tilde{\vdash} \tilde{\eta}$) when $\hat{\eta} \neq \hat{ff}$ and $\hat{a} \hat{\vdash} \hat{\eta}$, but $\hat{a} \not\tilde{\vdash} \check{\eta}$. Given an abstract condition $\tilde{\eta}$, $\hat{a}, \hat{b} \in \hat{\mathcal{C}}$ and $\check{a} \in \check{\mathcal{C}}$, an abstract conditional tuple is either a triple $\tilde{\eta} \rightarrow \langle \hat{a}, \hat{b} \rangle^m$, such that $\hat{a} \tilde{\vdash} \tilde{\eta}$, or a construct of the form $stutt(\check{a})^m$, where $m \in \{0, +\infty\}$ states how many times the corresponding tuples appear consecutively in the sequence. Given a (concrete) conditional tuple t , we define its abstraction $\alpha(t)$ as

$$\begin{array}{l}
 \alpha((\eta^+, \eta^-) \rightarrow \langle a, b \rangle) = (\tau^+(\eta^+), \tau^-(\eta^-)) \rightarrow \langle \tau^+(a), \tau^-(b) \rangle^1 \\
 \alpha(stutt(C)) = stutt(\tau^-(C))^1.
 \end{array}$$

Now, an abstract conditional reactive sequence is a sequence of *different* abstract tuples $\tilde{t}_1 \dots \tilde{t}_m \dots$, maybe ended with \square . The natural number associated to each abstract conditional tuple is needed to keep synchronization among processes due to the particularly strong synchronization properties of the language, as already noticed in Alpuente *et al.* (2005).

The abstraction $\alpha(s)$ of a sequence of conditional tuples s is defined by structural induction on the form of its tuples. It collapses all the computation steps (conditional

tuples) that, after abstraction, coincide. Formally, $\alpha(\epsilon) = \epsilon$, $\alpha(\square) = \square$ and

$$\alpha(t \cdot r) := \begin{cases} \tilde{\eta} \rightarrow \langle \hat{a}, \hat{b} \rangle^{m+1} \cdot \tilde{r} & \text{if } \alpha(t) = \tilde{\eta} \rightarrow \langle \hat{a}, \hat{b} \rangle^1, \alpha(r) = \tilde{\eta} \rightarrow \langle \hat{a}, \hat{b} \rangle^m \cdot \tilde{r} \\ stutt(\check{a})^{m+1} \cdot \tilde{r} & \text{if } \alpha(t) = stutt(\check{a})^1, \alpha(r) = stutt(\check{a})^m \cdot \tilde{r} \\ \alpha(t) \cdot \alpha(r) & \text{otherwise} \end{cases}$$

We extend this definition to sets of conditional sequences in the natural way.

We denote by \mathbf{A} the domain $\alpha(\mathbb{M})$ of the sets of abstract conditional reactive sequences. By adjunction we derive the concretization function γ such that

$$(\mathbb{M}, \sqsubseteq, \sqcup, \sqcap, \perp, \top) \xleftrightarrow[\alpha]{\gamma} (\mathbf{A}, \leq, \bigvee, \bigwedge, \perp, \top)$$

where $a \leq a' \iff \gamma(a) \sqsubseteq \gamma(a')$.

This abstraction can be systematically lifted to the domain of interpretations: $\mathbb{I} \xleftrightarrow[\alpha]{\gamma} [\mathbf{MGC} \rightarrow \mathbf{A}]$ so that we can derive the optimal abstraction of $\mathcal{D}[D]$ simply as $\mathcal{D}^\alpha[D] := \alpha \circ \mathcal{D}[D] \circ \gamma$. The abstract interpretation theory ensures that $\mathcal{F}^\alpha[D] := \mathcal{D}^\alpha[D] \uparrow \omega$ is the best correct approximation of $\mathcal{F}[D]$.

It turns out that $\mathcal{D}^\alpha[D]_{\mathcal{F}^\alpha} = \lambda p(x). \bigvee_{p(x) : -A \in D} \mathcal{A}^\alpha[A]_{\mathcal{F}^\alpha}$, where $\mathcal{A}^\alpha[\cdot]_{\mathcal{F}^\alpha}$ is defined by structural induction on the syntax in a similar way as the concrete version. Given the similarity to the concrete case, in the following we describe only two cases to illustrate the use of the upper- and lower-approximations (for full details consult (Comini et al. 2011)). The semantics for the tell agent just applies the abstraction to the only concrete sequence, thus: $\mathcal{A}^\alpha[\text{tell}(c)]_{\mathcal{F}^\alpha} = \{(\hat{t}, \check{f}) \rightarrow \langle \hat{t}, \tau^+(c) \rangle^1 \cdot \square\}$. For the now semantics, we only show the general case when the condition holds, and the general case when it does not hold:

$$\begin{aligned} \mathcal{A}^\alpha[\text{now}(d) \text{ then } A \text{ else } B]_{\mathcal{F}^\alpha} = & \\ & \{(d \hat{\times} \hat{\eta}, \check{\eta}) \rightarrow \langle d \hat{\times} \hat{a}, d \hat{\times} \hat{b} \rangle^n \cdot (d \tilde{\circ} \check{s}) \mid (\hat{\eta}, \check{\eta}) \rightarrow \langle \hat{a}, \hat{b} \rangle^n \cdot \check{s} \in \mathcal{A}^\alpha[A]_{\mathcal{F}^\alpha}, d \hat{\times} \hat{a} \tilde{\triangleright} (d \hat{\times} \hat{\eta}, \check{\eta})\} \\ & \vee \dots \vee \\ & \{(\hat{\eta}, d \check{\times} \check{\eta}) \rightarrow \langle \hat{a}, \hat{b} \rangle^1 \cdot (\hat{\eta}, \check{\eta}) \rightarrow \langle \hat{a}, \hat{b} \rangle^n \cdot \check{s} \mid (\hat{\eta}, \check{\eta}) \rightarrow \langle \hat{a}, \hat{b} \rangle^{n+1} \cdot \check{s} \in \mathcal{A}^\alpha[B]_{\mathcal{F}^\alpha}, \hat{a} \tilde{\triangleright} (\hat{\eta}, d \check{\times} \check{\eta})\} \\ & \vee \dots \end{aligned}$$

the $\tilde{\circ}$ operator is the abstract counterpart of the concrete version.

5 Abstract diagnosis of timed concurrent constraint programs

Now, following the ideas of Comini et al. (1999), we define the abstract diagnosis of *tccp*. The framework of abstract diagnosis (Comini et al. 1999) comes from the idea of considering the abstract versions of Park's Induction Principle.⁵ It can be considered as an extension of declarative debugging since there are instances of the framework that deliver the same results. In the general case, diagnosing w.r.t. *abstract* properties relieves the user from having to specify in excessive detail the program behavior (which could be more error-prone than the coding itself).

Let us now introduce the workset of abstract diagnosis. Having chosen a property of the computation α of interest (an instance of the abstraction scheme of Section 4),

⁵ A concept of formal verification that is undecidable in general.

given a set of declarations D and $\mathcal{S}^\alpha \in \mathbf{A}$, which is the specification of the intended behavior of D w.r.t. the property α , we say that

1. D is (abstractly) *partially correct* w.r.t. \mathcal{S}^α if $\alpha(\mathcal{F}[[D]]) \leq \mathcal{S}^\alpha$.
2. D is (abstractly) *complete* w.r.t. \mathcal{S}^α if $\mathcal{S}^\alpha \leq \alpha(\mathcal{F}[[D]])$.
3. D is *totally correct* w.r.t. \mathcal{S}^α , if it is partially correct and complete.

In this setting, the user can only reason in terms of the properties of the expected concrete semantics without being concerned with (approximate) abstract computations. The *diagnosis* determines the “originating” symptoms and, in the case of incorrectness, the relevant process declaration in the program. This is captured by the definitions of *abstractly incorrect process declaration* and *abstract uncovered element*:

Definition 5

Let D be a set of declarations, R a process declaration and $\{e\}, \mathcal{S}^\alpha \in \mathbf{A}$.

R is *abstractly incorrect* w.r.t. \mathcal{S}^α if $\mathcal{D}^\alpha[\{R\}]_{\mathcal{S}^\alpha} \not\leq \mathcal{S}^\alpha$.

e is an *uncovered element* w.r.t. \mathcal{S}^α if $\{e\} \leq \mathcal{S}^\alpha$ and $\{e\} \wedge \mathcal{D}^\alpha[D]_{\mathcal{S}^\alpha} = \perp$.

Informally, R is abstractly incorrect if it derives a wrong abstract element from the intended semantics. e is uncovered if the process declarations cannot derive it from the intended semantics.

It is worth noting that the notions of correctness and completeness are defined in terms of $\alpha(\mathcal{F}[[D]])$, i.e., in terms of abstraction of the concrete semantics. The abstract version of algorithmic debugging (Shapiro 1982), which is based on symptoms (i.e., deviations between $\alpha(\mathcal{F}[[D]])$ and \mathcal{S}^α), requires the construction of $\alpha(\mathcal{F}[[D]])$ and therefore a fixpoint computation. In contrast, the notions of abstractly incorrect process declarations and abstract uncovered elements are defined in terms of *just one* application of $\mathcal{D}^\alpha[D]$ to \mathcal{S}^α . The issue of the precision of the abstract semantics is specially relevant in establishing the relation between the two concepts (i.e., the relation between abstractly incorrect process declarations and abstract uncovered elements on one side, and abstract partial correctness and completeness, on the other side).⁶

Theorem 1

1. If there are no abstractly incorrect process declarations in D , then D is partially correct w.r.t. \mathcal{S}^α .
2. Let D be partially correct w.r.t. \mathcal{S}^α . If D has abstract uncovered elements then D is not complete.

When applying the diagnosis w.r.t. approximate properties, the results may be weaker than those that can be achieved on concrete domains just because of approximation. Abstract incorrect process declarations are in general just a warning about a possible source of errors. Because of the approximation, it can happen that a (concretely) correct declaration is abstractly incorrect. However, as shown by the following theorem, all concrete errors are detected, as they lead to an abstract incorrectness or abstract uncovered.

⁶ Proofs are available at <http://www.dimi.uniud.it/comini/Papers>.

Theorem 2

Let r be a process declaration and \mathcal{S} a concrete specification.

1. If $\mathcal{D}[\{r\}]_{\mathcal{S}} \not\sqsubseteq \mathcal{S}$ and $\alpha(\mathcal{D}[\{r\}]_{\mathcal{S}}) \not\leq \alpha(\mathcal{S})$ then r is abstractly incorrect w.r.t. $\alpha(\mathcal{S})$.
2. If there exists an abstract uncovered element a w.r.t. $\alpha(\mathcal{S})$, such that $\gamma(a) \sqsubseteq \mathcal{S}$ and $\gamma(\perp) = \perp$, then there exists a concrete uncovered element e w.r.t. \mathcal{S} (i.e., $e \sqsubseteq \mathcal{S}$ and $e \sqcap \mathcal{D}[D]_{\mathcal{S}} = \perp$).

It is particularly useful for applications the fact that our proposal can be used with partial specifications and also with partial programs. Obviously, one cannot detect errors in process declarations involving processes which have not been specified, but for the process declarations that involve processes that have a specification, the check can be made, even if the whole program has not been written yet. This includes the possibility of applying our “local” method to all parts of a program not involving constructs which we cannot handle (yet). With other “global” approaches such programs could not be checked at all.

It is worthy to note that, even for a noetherian abstract constraint system $\hat{\mathbf{C}}$, the domain of abstract sequences defined above is not—in general—noetherian, due to the use of the index in each tuple (we cannot get rid of it since it is needed to keep synchronization among parallel processes). This means that our current proposal cannot be used for static program analysis, unless we resort to use widening operators. However (for noetherian abstract constraint systems) our abstract diagnosis is effective since specifications have to be abstractions of some concrete semantics and, since the store evolves monotonically, it holds that the number of conditional tuples that can appear in an abstract sequence is, thus, finite.

5.1 Examples of application of the framework

Let us now show two illustrative examples of the approach. The first example shows the new ability of our approach: that of dealing with the constructors that introduce the nonmonotonic behavior of the system, in particular the now agent.

Example 1

We model a (simplified) *time-out*(n) process that checks for, at most, n times units if the system emits a signal telling that the process evolves normally (*system = ok*). When the signal arrives, the system emits the fact that there is no alert (*alert = no*).⁷ Let d_0 , d_n , d_{action} be the following declarations:

$$\begin{aligned} \text{time-out}(0) &: - \text{now}(\text{system} = \text{ok}) \text{ then } \text{action} \text{ else } (\text{ask}(tt) \rightarrow \text{time-out}(0)) \\ \text{time-out}(n) &: - \text{now}(\text{system} = \text{ok}) \text{ then } \text{action} \text{ else } (\text{ask}(tt) \rightarrow \text{time-out}(n-1)) \\ \text{action} &: - \text{tell}(\text{alert} = \text{no}) \end{aligned}$$

When the time limit is reached (declaration d_0), the system should set the signal *alert* to *yes* ($\text{tell}(\text{alert} = \text{no})$). However, we have introduced an error in the program, calling the process recursively instead: *time-out*(0).

⁷ The classical timeout would restart the countdown by recursively calling *time-out*(n).

Due to the simplicity of the constraint system, the abstract domain coincide with the concrete one, and the two external functions are the $\hat{\otimes}$ and $\hat{\oplus}$ operators.

Let us now consider the following specification. For d_0 we expect that, if the *ok* signal is present, then it ends with an *alert = no* signal, otherwise an alert should be emitted. This is represented by two possible sequences, one with a condition where *system = ok*, and a second one when *system = ok* is absent (this is a sequence that reasons with the *absence* of information).

$$\begin{aligned} \mathcal{S}^\alpha(\text{time-out}(0)) = & \{ (\text{system} = \text{ok}, \check{f}) \rightarrow \langle \text{system} = \text{ok}, \text{system} = \text{ok} \rangle^1 \cdot \\ & (\hat{t}, \check{f}) \rightarrow \langle \text{system} = \text{ok}, \text{system} = \text{ok} \hat{\otimes} \text{alert} = \text{no} \rangle^1 \cdot \square \} \\ & \cup \{ (\hat{t}, \{ \text{system} = \text{ok} \}) \rightarrow \langle \hat{t}, \hat{t} \rangle^1 \cdot (\hat{t}, \check{f}) \rightarrow \langle \hat{t}, \text{alert} = \text{yes} \rangle^1 \cdot \square \} \end{aligned}$$

The specification for d_n is similar, but we add n sequences, since we have the possibility that the signal arrives at each time instant before n .

$$\begin{aligned} \mathcal{S}^\alpha(\text{time-out}(n)) = & \{ (\hat{t}, \{ \text{system} = \text{ok} \}) \rightarrow \langle \hat{t}, \hat{t} \rangle^m \cdot \\ & (\text{system} = \text{ok}, \check{f}) \rightarrow \langle \text{system} = \text{ok}, \text{system} = \text{ok} \rangle^1 \cdot \\ & (\hat{t}, \check{f}) \rightarrow \langle \text{system} = \text{ok}, \text{system} = \text{ok} \hat{\otimes} \text{alert} = \text{no} \rangle^1 \cdot \square \mid 0 \leq m < n \} \\ & \cup \{ (\hat{t}, \{ \text{system} = \text{ok} \}) \rightarrow \langle \hat{t}, \hat{t} \rangle^{n+1} \cdot (\hat{t}, \check{f}) \rightarrow \langle \hat{t}, \text{alert} = \text{yes} \rangle^1 \cdot \square \} \\ \mathcal{S}^\alpha(\text{action}) = & \{ (\hat{t}, \check{f}) \rightarrow \langle \hat{t}, \text{alert} = \text{no} \rangle^1 \cdot \square \} \end{aligned}$$

Now, when we compute $\mathcal{D}^\alpha[\{d_0\}]_{\mathcal{S}^\alpha}$ we have:

$$\begin{aligned} & \{ (\text{system} = \text{ok}, \check{f}) \rightarrow \langle \text{system} = \text{ok}, \text{system} = \text{ok} \rangle^1 \cdot \\ & (\hat{t}, \check{f}) \rightarrow \langle \text{system} = \text{ok}, \text{system} = \text{ok} \hat{\otimes} \text{alert} = \text{no} \rangle^1 \cdot \square \} \\ & \cup \{ (\hat{t}, \{ \text{system} = \text{ok} \}) \rightarrow \langle \hat{t}, \hat{t} \rangle^1 \cdot (\text{system} = \text{ok}, \check{f}) \rightarrow \langle \text{system} = \text{ok}, \text{system} = \text{ok} \rangle^1 \cdot \\ & (\hat{t}, \check{f}) \rightarrow \langle \text{system} = \text{ok}, \text{system} = \text{ok} \hat{\otimes} \text{alert} = \text{no} \rangle^1 \cdot \square \} \\ & \cup \{ (\hat{t}, \{ \text{system} = \text{ok} \}) \rightarrow \langle \hat{t}, \hat{t} \rangle^2 \cdot (\hat{t}, \check{f}) \rightarrow \langle \hat{t}, \text{alert} = \text{no} \rangle^1 \cdot \square \} \end{aligned}$$

Due to the last sequence, $\mathcal{D}^\alpha[\{d_0\}]_{\mathcal{S}^\alpha} \not\leq \mathcal{S}^\alpha$, so we conclude that d_0 is (abstractly) incorrect. This is due to the recursive call in the else branch of the declaration. If we fix the program replacing d_0 by d'_0 where the recursive call is replaced by $\text{tell}(\text{alert} = \text{yes})$, then $\mathcal{D}^\alpha[\{d'_0\}]_{\mathcal{S}^\alpha} \leq \mathcal{S}^\alpha$, thus d'_0 is abstractly correct.

In Falaschi *et al.* (2007) it was studied an example where a *control* process checks whether a *failure* signal arrives to the system. The most important point that differs from the *timeout* example is that, in the *control* case, someone has to explicitly tell the system that an error has occurred. Instead, in the *timeout* example, the system is able to act (and maybe recover) when it detects that something that should have happened, had not happened. In other words, the *control* example does not handle *absence* of information, since nonmonotonic operators are not considered there. We have implemented the example in *tcp* and we have checked that the same results can be achieved in our framework if we apply the same abstraction they use (a *depth(k)* abstraction).

The second example we show illustrates how one can work with the abstraction of the constraint system, and also how we can take advantage of our abstract domain.

Example 2

Let us consider a system with a single declaration and the abstraction of the constraint system that abstracts integer variables to a (simplified) interval-based domain with abstract values $\{\top, \text{pos}_x, \text{neg}_x, x > 10, x \leq 10, \perp\}$.

$$p(x) :- \text{now}(x > 10) \text{ then } \exists x' (\text{tell}(x = \lfloor _ | x' \rfloor) \parallel \text{tell}(x' = \lfloor x + 1 | _ \rfloor) \parallel p(x')) \\ \text{else } \exists x'' (\text{tell}(x = \lfloor _ | x'' \rfloor) \parallel \text{tell}(x'' = \lfloor x - 1 | _ \rfloor) \parallel p(x''))$$

Due to the monotonicity of the store, we have to use streams (written in a list-fashion way) to model the *imperative-style* variables (de Boer et al. 2000). In this way, variable x in the program above is a stream that is updated with different values during the execution. Following this idea, the abstraction for concrete streams is defined as the (abstracted) last instantiated value in the stream. The concretization of one stream is defined as all the concrete streams whose last value is a concretization of the abstract one. We write a dot on a predicate symbol (e.g. $\dot{=}$) to denote that we want to check it for the last instantiated value of a stream.

We define the following intended specification to specify that, (a) if the parameter is greater than 10, then the last value of the stream (written \dot{x}) will always be greater than 10; (b) if the parameter is negative, then the value is always negative

$$\mathcal{S}^\alpha(p(x_1)) = \{(x_1 > 10, \check{f}) \rightarrow \langle x > 10, x > 10 \rangle^{+\infty}\} \cup \{(\text{neg}_{\dot{x}}, \check{f}) \rightarrow \langle \text{neg}_{\dot{x}}, \text{neg}_{\dot{x}} \rangle^{+\infty}\}$$

The two abstract sequences represent infinite computations thanks to the $+\infty$ index in the last tuple. In other words, finite specifications that represent infinite computations can be considered and effectively handled. In fact, we can compute $\mathcal{D}^\alpha[\{d\}]^{\mathcal{S}^\alpha}$:

$$\begin{aligned} & \{ \{ (\text{neg}_{\dot{x}}, \check{f}) \rightarrow \langle \text{pos}_{\dot{x}}, \text{pos}_{\dot{x}} \rangle^1 \cdot (\text{pos}_{\dot{x}} \tilde{\circ} (\dot{x} > 10, \check{f}) \rightarrow \langle \dot{x} > 10, \dot{x} > 10 \rangle^{+\infty}) \} \\ & \cup \{ (\text{neg}_{\dot{x}}, \check{f}) \rightarrow \langle \text{neg}_{\dot{x}}, \text{neg}_{\dot{x}} \rangle^1 \cdot (\text{neg}_{\dot{x}}, \check{f}) \rightarrow \langle \text{neg}_{\dot{x}}, \text{neg}_{\dot{x}} \rangle^{+\infty} \} \} \\ = & \\ & \{ \{ (\text{neg}_{\dot{x}}, \check{f}) \rightarrow \langle \text{pos}_{\dot{x}}, \text{pos}_{\dot{x}} \rangle^1 \cdot \overbrace{(\text{pos}_{\dot{x}} \hat{\circ} \dot{x} > 10, \check{f})}^{\text{pos}_{\dot{x}}} \rightarrow \langle \text{pos}_{\dot{x}} \hat{\circ} \dot{x} > 10, \text{pos}_{\dot{x}} \hat{\circ} \dot{x} > 10 \rangle^{+\infty} \} \\ & \cup \{ (\text{neg}_{\dot{x}}, \check{f}) \rightarrow \langle \text{neg}_{\dot{x}}, \text{neg}_{\dot{x}} \rangle^1 \cdot (\text{neg}_{\dot{x}}, \check{f}) \rightarrow \langle \text{neg}_{\dot{x}}, \text{neg}_{\dot{x}} \rangle^{+\infty} \} \} \\ = & \\ & \{ \{ (\text{pos}_{\dot{x}}, \check{f}) \rightarrow \langle \text{pos}_{\dot{x}}, \text{pos}_{\dot{x}} \rangle^{+\infty} \} \cup \{ (\text{neg}_{\dot{x}}, \check{f}) \rightarrow \langle \text{neg}_{\dot{x}}, \text{neg}_{\dot{x}} \rangle^{+\infty} \} \} \end{aligned}$$

The third equality holds because $\text{pos}_{\dot{x}}$ entails $x > 10$, so the merge of the two constraints will be equal to $\text{pos}_{\dot{x}}$.

Since $\mathcal{D}^\alpha[\{d\}]^{\mathcal{S}^\alpha} \not\leq \mathcal{S}^\alpha$ we can conclude that d is an incorrect declaration w.r.t. \mathcal{S}^α . In addition, we can notice that \mathcal{S}^α contain an uncovered element that is a sequence that cannot be derived by the semantics operator.

6 Related work

A top-down (big-step) denotational semantics for tcp is defined in de Boer et al. (2000) for terminating computations. In that work, a terminating computation is both, a computation that reaches a point in which no agents are pending to be executed, and also a computation that suspends since there is no enough information in the store to make the choice agents evolve. Our semantics is a

bottom-up (small-step) denotational semantics that models infinite computations, and also distinguishes the two kinds of terminating computations aforementioned. Conceptually, a suspended computation has not completely finished its execution, and, in some cases, it could be a symptom of a system error. Thus, the new semantics is well suited to handle, not only functional systems (where an input-output semantics makes sense), but also reactive systems.

In Falaschi *et al.* (2007), a first approach to the declarative debugging of a ccp language is presented. However, it does not cover the particular extra difficulty of the nonmonotonicity, common to all timed concurrent constraint languages. As we have said, this ability is crucial to model specific behaviors of reactive systems, such as timeouts or preemption actions. This is the main reason why our abstract (and concrete) semantics are significantly different from Falaschi *et al.* (2007) and from formalizations for other declarative languages.

The idea of using two different mechanisms for dealing with positive and negative information in our abstraction scheme is inspired by Alpuente *et al.* (2005). There, a framework for the abstract model checking of tcp programs based on a source-to-source transformation is defined. In particular, it is defined a transformation from a tcp program P into a tcp program \bar{P} that represents a correct abstraction of the original one (in the sense that the semantics of P are included in the semantics of \bar{P}). Instead, we define an abstract semantics for the language. The upper- and lower-approximated versions of the entailment relation are used to keep \bar{P} correct, but also precise enough.

7 Conclusion and future work

We have presented a new compact, bottom-up semantics for the tcp language which is correct and fully abstract w.r.t. the behavior of the language. This semantics is well suited for debugging and verification purposes in the context of reactive systems. The idea of using conditions to have a correct bottom-up semantics can be also applied to other nonmonotonic languages such as, for example, ntcc in the ccp paradigm (Palamidessi and Valencia 2001) or Linda in the imperative (coordination) paradigm (Gelernter 1985).

Then, an abstract semantics that is able to specify (a kind of) infinite computations is presented. It is based on the abstraction of computation sequences by using two functions that satisfy some properties in order to guarantee correctness. All our examples satisfy those conditions. The abstract semantics keeps the synchronization among parallel computations, which is a particular difficulty of the tcp language. As already noticed in Alpuente *et al.* (2005), the loss of synchronization in other ccp languages just implies a loss of precision, but in the case of tcp, due to the maximal parallelism, it would imply a loss of correctness.

Finally, we have adapted the abstract diagnosis approach to the tcp language employing the new semantics as basis. We have presented two illustrative examples to show the new features of our approach w.r.t. other paradigms.

As future work, we intend to work on abstractions of our semantics to domains of temporal logic formulas, to be able to specify safety and/or liveness properties,

and to compare its models w.r.t. the program semantics. Another interesting aspect is to study if a general framework for the proposed methodology can be defined to apply it to other languages.

References

- ALPUENTE, M., COMINI, M., ESCOBAR, S., FALASCHI, M. AND LUCAS, S. 2003. Abstract Diagnosis of Functional Programs. In *Logic Based Program Synthesis and Transformation—12th International Workshop, LOPSTR 2002, Revised Selected Papers*, M. Leuschel, Ed. Lecture Notes in Computer Science, vol. 2664. Springer-Verlag, Berlin, 1–16.
- ALPUENTE, M., GALLARDO, M., PIMENTEL, E. AND VILLANUEVA, A. 2005. A semantic framework for the abstract model checking of tccp programs. *Theoretical Computer Science* 346(1), 58–95.
- BACCI, G. AND COMINI, M. 2011. Abstract diagnosis of first order functional logic programs. In *Logic-based Program Synthesis and Transformation, 20th International Symposium*, M. Alpuente, Ed. Lecture Notes in Computer Science, vol. 6564. Springer-Verlag, Berlin, 208–226.
- COMINI, M., LEVI, G., MEO, M. C. AND VITIELLO, G. 1999. Abstract diagnosis. *Journal of Logic Programming* 39(1–3), 43–93.
- COMINI, M., TITOLO, L. AND VILLANUEVA, A. 2011. A Compact Goal-Independent Bottom-Up Fixpoint Modeling the Behavior of tccp. Tech. Rep. DIMI-UD/01/2011/RR, Dipartimento di Matematica e Informatica, U. di Udine. <http://www.dimi.uniud.it/comini/Papers/>.
- COUSOT, P. AND COUSOT, R. 1979. Systematic Design of Program Analysis Frameworks. In *Proceedings of the 6th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, San Antonio, Texas, January 29–31. ACM Press, New York, NY, USA, 269–282.
- DE BOER, F. S., GABBRIELLI, M. AND MEO, M. C. 2000. A timed concurrent constraint language. *Information and Computation* 161(1), 45–83.
- FALASCHI, M., OLARTE, C., PALAMIDESSI, C. AND VALENCIA, F. 2007. Declarative diagnosis of temporal concurrent constraint programs. In *Proceedings of the 23rd International Conference on Logic Programming (ICLP'07)*. Lecture Notes in Computer Science, vol. 4670. Springer-Verlag, 271–285.
- GELERNTER, D. 1985. Generative Communication in Linda. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 7(1), 80–113.
- PALAMIDESSI, C. AND VALENCIA, F. D. 2001. A Temporal Concurrent Constraint Programming Calculus. In *Proceedings of the 7th International Conference on Principles and Practice of Constraint Programming (CP'01)*. Lecture Notes in Computer Science, vol. 2239. Springer, 302–316.
- SARASWAT, V. A. 1993. *Concurrent Constraint Programming*. The MIT, Cambridge, MA.
- SHAPIRO, E. Y. 1982. Algorithmic Program Debugging. In *Proceedings of Ninth Annual ACM Symp. on Principles of Programming Languages*. ACM Press, New York, NY, USA, 412–531.