

*Modeling and Reasoning in Event Calculus using Goal-Directed Constraint Answer Set Programming *[†]*

JOAQUÍN ARIAS

CETINIA, Universidad Rey Juan Carlos, Madrid, Spain
(e-mail: joaquin.arias@urjc.es)

MANUEL CARRO

IMDEA Software Institute, Madrid, Spain
Universidad Politécnica de Madrid, Madrid, Spain
(e-mails: manuel.carro@imdea.org, manuel.carro@upm.es)

ZHUO CHEN

University of Texas at Dallas, Richardson, USA
(e-mail: zhuo.chen@utdallas.edu)

GOPAL GUPTA

University of Texas at Dallas, Richardson, USA
(e-mail: gupta@utdallas.edu)

submitted 21 April 2021; revised 29 June 2021; accepted 08 July 2021

Abstract

Automated commonsense reasoning (CR) is essential for building human-like AI systems featuring, for example, explainable AI. Event calculus (EC) is a family of formalisms that model CR with a sound, logical basis. Previous attempts to mechanize reasoning using EC faced difficulties in the treatment of the continuous change in dense domains (e.g. time and other physical quantities), constraints among variables, default negation, and the uniform application of different inference methods, among others. We propose the use of s(CASP), a query-driven, top-down execution model for Predicate Answer Set Programming with Constraints, to model and reason using EC. We show how EC scenarios can be naturally and directly encoded in s(CASP) and how it enables deductive and abductive reasoning tasks in domains featuring constraints involving both dense time and dense fluents.

KEYWORDS: event calculus, constraints, ASP, commonsense reasoning

1 Introduction

The ability to model continuous characteristics of the world is essential for commonsense reasoning (**CR**) in many domains that require dealing with continuous change: time, the

* This paper is an extended version of the work by Arias *et al.* (2019).

[†] Work partially supported by EIT Digital, MICINN projects RTI2018-095390-B-C33 InEDGEMobility (MCIU/AEI/FEDER, UE), PID2019-108528RB-C21 ProCode, Comunidad de Madrid project S2018/TCS-4339 BLOQUES-CM co-funded by EIE Funds of the European Union, US NSF Grants IIS 1718945, IIS 1910131, IIP 1916206.

height of a falling object, the gas level of a car, the water level in a sink, etc. Event calculus (**EC**) is a formalism based on many-sorted predicate logic (Kowalski and Sergot 1989; Mueller 2014) that can represent continuous change and capture the commonsense law of inertia, whose modeling is a pervasive problem in CR. In EC, time-dependent properties and events are seen as objects and reasoning is performed on the truth values of properties and the occurrences of events at a point in time.

Answer set programming (**ASP**) is a logic programming paradigm that was initially proposed by Marek and Truszczyński (1999) and Lifschitz (1999) to realize non-monotonic reasoning. ASP has been used by Lee and Palla (2012; 2020) to model the EC. Classical implementations of ASP are limited to variables ranging over discrete and bound domains and use grounding and SAT solving to find out models (called *answer sets*) of ASP programs. However, reasoning on models of the real world often needs variables ranging over dense domains (domains that are continuous, such as \mathbb{R} , or that are not continuous but have an infinite number of elements in any bound, non-singleton interval, such as \mathbb{Q}). Dense domains are necessary to accurately represent the properties of some physical quantities, such as time, weight, space, etc.

This paper presents an approach to modeling EC using the s(CASP) system by Arias et al. (2018) as the underlying reasoning infrastructure. The s(CASP) system is an implementation of constraint answer set programming over first-order predicates which combines ASP and constraints. It features predicates, constraints among non-ground variables, uninterpreted functions, and, most importantly, a top-down, query-driven execution strategy. These features make it possible to return answers with non-ground variables, possibly including constraints among them, and to compute partial models by returning only the fragment of a stable model that is necessary to support the answer to a given query. Thanks to its interface with constraint solvers, sound non-monotonic reasoning with constraints is possible. This approach achieves more conciseness and expressiveness, in the sense of being able to succinctly express complex computations and reasoning tasks, than other related approaches. Dense domains can be faithfully modeled in s(CASP) as continuous quantities, while in other proposals such domains had to be discretized, as done by Mellarkod et al. (2008a) and Lee and Palla (2020), therefore losing precision or even soundness. Additionally, in our approach the amalgamation of ASP and constraints and its realization in s(CASP) is considerably more natural: under s(CASP), answer set programs are executed in a goal-directed manner so constraints encountered along the way are collected and solved dynamically as execution proceeds — this is very similar to the way in which Prolog was extended with constraints. The implementation of other ASP systems featuring constraints is considerably more complex.

In the rest of the paper we present s(CASP) and its unique capabilities together with a terse introduction to EC (Section 2), our approach to modeling EC with s(CASP) (Section 3), a quantitative and qualitative evaluation (Section 4), and, finally, related work and conclusions (Sections 5 and 6).

2 Background

ASP is a logic programming and modeling language that evaluates normal logic programs under the stable model semantics proposed by Gelfond and Lifschitz (1988). s(ASP), introduced by Marple et al. (2017), is a top-down, goal-driven ASP system that can evaluate ASP programs with function symbols (*functors*) **without** *grounding* them either

before or during execution. Grounding is a procedure that substitutes program variables with the possible values from their domain. For most classical ASP solvers, grounding is a necessary pre-processing phase. Grounding, however, requires program variables to be restricted to take values in a finite domain. As a result, traditional ASP solvers cannot be used to model continuous time or change.

2.1 *s(CASP)*

s(CASP), presented by Arias *et al.* (2018), extends *s(ASP)* by adding constraints, similarly to how CLP extends Prolog. Also, similarly to how *s(ASP)* can compute the stable model semantics with non-ground programs and return non-ground models including disequalities in the Herbrand domain, *s(CASP)* keeps constraints as relations among variables both during execution and in the answer sets.

Constraints have historically proved to be effective in improving both expressiveness (programs are shorter and easier to understand, as many computation details are taken care of by the underlying constraint solver) and efficiency in logic programming, as they can succinctly express properties of a solution and reduce the search space. As a result, *s(CASP)* is more expressive and faster than *s(ASP)*, while retaining the capability of executing non-ground predicate answer set programs.

2.1.1 Syntax and behavior

An *s(CASP)* program is a set of clauses of the following form:

$$\mathbf{a} :- \mathbf{c}_a, \mathbf{b}_1, \dots, \mathbf{b}_m, \text{not } \mathbf{b}_{m+1}, \dots, \text{not } \mathbf{b}_n.$$

where \mathbf{a} and $\mathbf{b}_1, \dots, \mathbf{b}_n$ are atoms. An atom is either a propositional variable or the expression $\mathbf{p}(\mathbf{t}_1, \dots, \mathbf{t}_n)$ if \mathbf{p} is an n -ary predicate symbol and $\mathbf{t}_1, \dots, \mathbf{t}_n$ are terms. A term is either a variable x_i or a function symbol \mathbf{f} of arity \mathbf{n} , denoted as \mathbf{f}/\mathbf{n} , applied to n terms, for example, $\mathbf{f}(\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_n)$, where each \mathbf{t}_i is in turn a term. A function symbol of arity 0 is called a constant. Program variables are usually written starting with an uppercase letter, while function and predicate symbols start with a lowercase letter. Numerical constants are written solely with digits.¹ Therefore, *s(CASP)* accepts terms with the same conventions as Prolog: $\mathbf{f}(\mathbf{a}, \mathbf{b})$ is a term, and so are $\mathbf{f}(\mathbf{g}(\mathbf{X}), \mathbf{Y})$ and $[\mathbf{f}(\mathbf{a}) | \text{Rest}]$ (to denote a list with head $\mathbf{f}(\mathbf{a})$ and tail *Rest*).

A note on terminology *ASP literature often uses the term constraint to denote constructions such as $:- \mathbf{p}, \mathbf{q}$, that is, rules without head. They express that the conjunction of atoms $\mathbf{p} \wedge \mathbf{q}$ cannot be true: either \mathbf{p} , \mathbf{q} , or both, have to be false in any stable model. Our programs use as well a designated class of predicates that restrict the domains of variables (Marriott and Stuckey 1998; Marriott *et al.* 2006) and whose semantics comes from the structure and domain of an underlying constraint system. In Constraint Logic Programming (CLP), these predicates are called “constraints” as well. To avoid the ambiguity that may arise from using the same name for constraints appearing among (free) variables during program execution and in the final models and for rules*

¹ There are additional syntactical conventions to separate variables and non-variables that are of no interest here.

without heads, and following Mellarkod et al. (2008b), we will refer to headless rules as denials.

c_a is a simple constraint or a conjunction of constraints: an expression establishing relations among variables in some constraint system, as described by Marriott and Stuckey (1998). Similar to CLP, s(CASP) is parametrized by the constraint system, from which it inherits its semantics. Since the execution of an s(CASP) program needs negating constraints (Section 2.1.3), we require that this can be done in the constraint system by means of a finite disjunction of basic constraints (Stuckey 1991; Dovier et al. 2000).

At least one of a , b_i , $\text{not } b_i$, or c_a must be present. When the head a is not present it is supposed to be substituted by the head *false*. The rules have then the form

$$:- c_a, b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n.$$

(and, as mentioned before, we call this rule a *denial*) and their interpretation is that the conjunction of the constraints and goals has to be false, so at least one constraint or goal has to be false.

The execution of an s(CASP) program starts with a *query* of the form

$$?- c_a, b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n.$$

The s(SCASP) answers to a query are *partial* stable models where each one is a subset of a stable model that satisfies the constraints, makes non-negated atoms true, makes the negated atoms non-provable, and, in addition, includes only atoms that are relevant to support the query. Additionally, for each partial stable model s(CASP) can return on backtracking the justification tree and the bindings for the free variables of the query that correspond to the most general unifier (*mgu*) of a successful top-down derivation consistent with this stable model.

An atom can have the form -r (i.e. have a hyphen as its first character). In that case it is assumed to express the *classical negation* of atom r . Rules with head -r , to express when r is false, can be part of the program. To ensure soundness, a denial

$$:- r, \text{-r}.$$

is automatically added to guarantee that atom r and its classical negation -r are not both simultaneously true in any model. Other than that, -r is not treated specially by s(CASP). The construct $\text{not } \text{-r}$ is allowed and rules with -r in their head or body are subject to *dualization* (Section 2.1.3).

Default negation $\text{not } r$ differs from classical negation -r in that $\text{not } r$ succeeds when it cannot be proven from the program that r holds, while -r succeeds if there is a rule that states how to deduce -r and this rule, together with the rest of the program, can be used to derive -r . The difference from the point of view of reasoning can be illustrated with a simple piece of commonsense knowledge: a bus may cross the railway tracks if no train is approaching. A possible rule using classical negation expresses would be:

$$\text{cross} :- \text{-train}.$$

It means the railway tracks can be crossed if we *explicitly* know (because there is a proof for it) that no train is approaching — for example, because there are sensors that send us information that ensures that there is definitely no train on the tracks within some safety range. The rule using default negation would be:

$$\text{cross} :- \text{not train}.$$

which means that we can cross the railway tracks if there is no evidence (because we cannot prove it) that a train is approaching — for example, we do not receive information of a train coming. But there may be no train coming, or the sensors may be faulty and not sending signals. That is why $\neg r$ is sometimes referred to as *strong* negation: it carries with it the meaning that there is a constructive proof that r is false. Therefore, lack of evidence is not a hard proof. Classical and default negation have two different meanings inside the language and are used to express very different CR scenarios.

When $\neg r$ and r are defined, the decision to invoke $\neg r$ or **not** r in the body a rule depends on what the programmer wants to express. There is a relation of containment between **not** r and $\neg r$, but it is clearer in the context of non-propositional atoms. Therefore, we defer its explanation to the end of Section 3.4, when we deal with the translations of the axioms of the Basic Event Calculus (BEC).

In *s(CASP)*, and unlike Prolog's negation as failure and ASP default negation, **not** $p(X)$ can return bindings for X on success, that is, bindings for which the call $p(X)$ would have failed. Constraints may be returned as well: for the program

```
1 p(a).
```

the query `?- not p(X).` would return the binding $X \neq a$ and the model $\{\text{not } p(X \mid \{X \neq a\})\}$, representing the set of *not* $p(X)$ such that the atom $p(X)$ can be proven only when $X \neq a$.² Note that X in the query appeared only in a negated atom, and did not need to be part of any non-negated atom. This is possible thanks to the use of constructive negation (Marple *et al.* 2017) and coinductive success (Gupta *et al.* 2007) in *s(ASP)*.

These are augmented in *s(CASP)* with the constraint processing capabilities presented in Arias *et al.* (2018), such that the program

```
1 p(X):- X > 0.
```

will return, for the same query as before, the model $\{\text{not } p(X \mid \{X \leq 0\})\}$.

s(CASP) uses a top-down, goal-driven execution procedure that implements an extension of the stable model semantics introduced by Gelfond and Lifschitz (1988) for non-ground programs. Default negation is solved against the *dual rules* of the program, which give a constructive definition of the negation of program predicates. The top-down algorithm does not need grounded programs and makes it possible as well to return partial stable models.

2.1.2 Overview of the execution procedure of *s(CASP)*

In the following sections we will present an abridged description of the top-down evaluation procedure used by *s(CASP)*, which, in a nutshell, is:

1. Rules expressing the constructive negation of the predicates in the original ASP program are synthesized (Section 2.1.3). We call this the *dual program*. Its mission is to provide a means to constructively determine the conditions and constraints under which calls to non-propositional predicates featuring variables would have failed: if

² Uniqueness of names is assumed for constants and function names: any two constants or functions with different names represent different objects.

we want to know when a rule such as $p(\mathbf{X}, \mathbf{Y}) :- q(\mathbf{X}), \text{not } r(\mathbf{Y})$ succeeds, the dual program computes the constraints on \mathbf{Y} under which the call $r(\mathbf{Y})$ would fail. This is an extension of the usual ASP semantics that is compatible with the case of programs that can be finitely grounded.³

2. The original program is checked for loops of the form $p :- q, \text{not } p$. and denials are generated for them.
3. The denials generated in point 2, together with any denials present in the original program, are collected in a predicate synthesized by the compiler that is invoked by adding an auxiliary goal to this predicate at the end of the query.
4. The union of the original program, the dual program, and the denials is handled by a top-down execution algorithm that implements the stable model semantics.

Item number 4 is specially relevant. The dual program (item 1) is synthesized by means of program transformations drawing from classical logic. However, its meaning differs from that of first-order logic. That is so because it is to be executed by a metainterpreter that does not implement the inference mechanisms of first-order logic, as it is designed to ensure that the semantics of answer set programs is respected. In particular, it treats specifically cyclic dependencies involving negation — see Section 2.1.4. For conciseness, we have not included in this paper the description of the execution algorithm, which can be found in Marple et al. (2012) and Marple et al. (2017).

Therefore, the soundness of $s(\text{CASP})$ (and its version without constraints, $s(\text{ASP})$) needs to be assessed taking into account the dual program, the generation of the denials and the evaluation algorithm as a whole. This was done by Marple et al. (2012) for the propositional case (which lays the bases of the whole procedure) and extended for the case of predicate logic, including arbitrary function symbols, by Marple et al. (2017). We will provide reasons supporting the soundness of the $s(\text{CASP})$ algorithm in the next sections.

2.1.3 Dual programs

We summarize here the synthesis of the dual of a logic program P : the completion procedure described by Clark (1978) is performed to generate a program $\text{Comp}(P)$, its rules are converted into an equivalent form with negated heads, and then De Morgan's laws are applied to generate separate clauses.

1. Following Ferraris et al. (2011, Section 2.1), first-order sentences are constructed for each clause by considering each i -th rule of predicate p

$$p :- c_a, b_1, \dots, b_m, \text{not } b_{m+1}, \dots, \text{not } b_n.$$

as a shorthand for $\forall \vec{x} \vec{y}_i (p_i(\vec{x}) \leftarrow B_i)$, where B_i corresponds to the conjunction $c_a \wedge b_1 \wedge \dots \wedge b_m \wedge \neg b_{m+1} \wedge \dots \wedge \neg b_n$ and \vec{y}_i are the variables appearing in B_i that do not appear in \vec{x} . The rationale for this transformation is that the grounding of an ASP program substitutes variables in the program clauses for all the constants

³ Note that, in the presence of function symbols and constraints on dense domains, this is in general not the case for $s(\text{CASP})$ programs.

in the program, and all the resulting clauses have to be satisfied. This is precisely what the universal quantifier expresses.

We will assume that clauses are normalized, that is, head unifications have been made explicit as goals in the bodies of the corresponding clauses and the heads contain only variable names. Also, in what follows we will not distinguish user predicates b_i from constraints c_a until the last step of the generation of the dual program; we will make the necessary distinction there.

2. All sentences corresponding to the same predicate name are conjoined together:

$$\begin{aligned} &\forall \vec{x} \vec{y}_i (p(\vec{x}) \leftarrow B_1) \wedge \\ &\quad \vdots \\ &\forall \vec{x} \vec{y}_k (p(\vec{x}) \leftarrow B_k). \end{aligned}$$

3. The bodies in the antecedent of the sentences are joined in a single body:

$$\forall \vec{x} \vec{y} (p(\vec{x}) \leftarrow B_1 \vee \dots \vee B_k).$$

The variables \vec{y} are updated to include all \vec{y}_i that appear in the different B_i and do not appear in \vec{x} .

4. The scope of the quantifiers is minimized to make further simplifications possible:

$$\forall \vec{x} (p(\vec{x}) \leftarrow \exists \vec{y} (B_1 \vee \dots \vee B_k)),$$

and then

$$\forall \vec{x} (p(\vec{x}) \leftarrow \exists \vec{y}_1 B_1 \vee \dots \vee \exists \vec{y}_k B_k).$$

Transformations 2 to 4 are valid in intuitionistic logic, and so they preserve the stable models of the original formulae, as mentioned by Ferraris *et al.* (2011, Section 6.1).

5. Implications are replaced by equivalences, to generate the Clark completion of the original predicate:

$$\forall \vec{x} (p(\vec{x}) \longleftrightarrow \exists \vec{y}_1 B_1 \vee \dots \vee \exists \vec{y}_k B_k).$$

This transformation is, in general, not model-preserving, except in the case of tight programs as presented by Erdem and Lifschitz (2003). For a program P , positive loops makes the Clark completion $Comp(P)$ under the classical first-order semantics be weaker than P under the stable model semantics: all stable models of P are classical models of $Comp(P)$, but not the other way around. Therefore, there may be classical models of $Comp(P)$ that are not stable models of P .

6. We create new predicate names to separate the bodies corresponding to the different original clauses:

$$\begin{aligned} &\forall \vec{x} (p(\vec{x}) \longleftrightarrow p_1(\vec{x}) \vee \dots \vee p_k(\vec{x})) \\ &\forall \vec{x} (p_i(\vec{x}) \longleftrightarrow \exists \vec{y}_i B_i). \end{aligned}$$

7. Their duals $\neg p/n$, $\neg p_i/n$ are:

$$\begin{aligned} &\forall \vec{x} (\neg p(\vec{x}) \longleftrightarrow \neg(p_1(\vec{x}) \vee \dots \vee p_k(\vec{x}))) \\ &\forall \vec{x} (\neg p_i(\vec{x}) \longleftrightarrow \neg \exists \vec{y}_i B_i). \end{aligned}$$

This is a semantically-preserving operation in the classical logic semantics, and so the models of the Clark completion remain untouched.

8. De Morgan's laws are applied to the first formula (the "entry point" of the negated predicate) and the existential quantifier is negated:

$$\begin{aligned} \forall \vec{x} (\neg p(\vec{x}) \longleftrightarrow \neg p_1(\vec{x}) \wedge \dots \wedge \neg p_k(\vec{x})) \\ \forall \vec{x} (\neg p_i(\vec{x}) \longleftrightarrow \forall \vec{y}_i \neg B_i). \end{aligned}$$

De Morgan's Law are of course semantics-preserving in classical logic, but also in the stable model semantics because the formulas $\neg(A \wedge B)$ and $\neg A \vee \neg B$ are strongly equivalent, as mentioned by Lifschitz et al. (2001, Proposition 5).

9. For each dual rule corresponding to each p_i , an auxiliary *negated* predicate corresponding to the *negated* body is synthesized:

$$\begin{aligned} \forall \vec{x} (\neg p(\vec{x}) \longleftrightarrow \neg p_1(\vec{x}) \wedge \dots \wedge \neg p_k(\vec{x})) \\ \forall \vec{x} (\neg p_i(\vec{x}) \longleftrightarrow \forall \vec{y}_i \neg p'_i(\vec{x}, \vec{y}_i)) \\ \forall \vec{x} \vec{y} (\neg p'_i(\vec{x}, \vec{y}) \longleftrightarrow \neg B_i). \end{aligned}$$

Let us remember that each B_i has the form $b_{i,1} \wedge \dots \wedge b_{i,m} \wedge \neg b_{i,m+1} \wedge \dots \wedge \neg b_{i,n}$. Therefore, the last sentence is

$$\forall \vec{x}, \vec{y} (\neg p'_i(\vec{x}, \vec{y}) \longleftrightarrow \neg(b_{i,1} \wedge \dots \wedge b_{i,m} \wedge \neg b_{i,m+1} \wedge \dots \wedge \neg b_{i,n})).$$

10. We apply De Morgan's Law to the last sentence in the previous point to obtain

$$\forall \vec{x}, \vec{y} (\neg p'_i(\vec{x}, \vec{y}) \longleftrightarrow \neg b_{i,1} \vee \dots \vee \neg b_{i,m} \vee b_{i,m+1} \vee \dots \vee b_{i,n}).$$

11. We revert the introduction of the equivalence. This transformation changes the models of program w.r.t. that of the Clark completion. However, programs under the stable semantics (and under Prolog semantics) have a clear notion of direction and the metainterpreter only uses the goal-driven direction "use *Body* to prove *Head*". The closed-world assumption, captured by Clark's completion, is already implicit in the top-down evaluation algorithm of s(CASP).

$$\begin{aligned} \forall \vec{x} (\neg p(\vec{x}) \leftarrow \neg p_1(\vec{x}) \wedge \dots \wedge \neg p_k(\vec{x})) \\ \forall \vec{x} (\neg p_i(\vec{x}) \leftarrow \forall \vec{y}_i \neg p'_i(\vec{x}, \vec{y}_i)) \\ \forall \vec{x} \vec{y} (\neg p'_i(\vec{x}, \vec{y}) \leftarrow \neg b_{i,1} \vee \dots \vee \neg b_{i,m} \vee b_{i,m+1} \vee \dots \vee b_{i,n}). \end{aligned}$$

12. Separate the disjunction in the body of the last sentence in different clauses:

$$\begin{aligned} \forall \vec{x} \vec{y} (\neg p'_i(\vec{x}, \vec{y}) \leftarrow \neg b_{i,1}) \\ \forall \vec{x} \vec{y} (\neg p'_i(\vec{x}, \vec{y}) \leftarrow \neg b_{i,2}) \\ \vdots \\ \forall \vec{x} \vec{y} (\neg p'_i(\vec{x}, \vec{y}) \leftarrow b_{i,n}). \end{aligned}$$

These clauses, together with the original program and the denials, are used by the metainterpreter to decide whether some atom belongs or not to a stable model of a program and to return the (minimal) support for that atom.

This provides a definition for $\neg p(\vec{x})$ via a clause with head $\neg p_i(\vec{x})$ for each original clause with head $p_i(\vec{x})$. The newly introduced negated atoms $\neg b_{i,1} \dots \neg b_{i,m}$ can fall into two categories: they are either negations of user predicates or negations of constraints. In the former case, the procedure just described generates a definition for the negation of user predicates. In the latter case, constraints are either head unifications created after normalizing the clauses, or they are constraints in a different domain. Both cases are treated similarly:

- If the negation of a constraint can be expressed as a finite disjunction of basic constraints (Dovier *et al.* 2000; Stuckey 1991), the compiler makes that expansion. In the simplest case, that disjunction is a single constraint: a linear constraint $E_1 < E_2$ is translated into $E_1 \geq E_2$. In other cases, it can be a “real” disjunction: the constraint $E_1 = E_2$ in CLP(Q) (linear constraints over the rationals), is negated by converting it into $E_1 < E_2 \vee E_1 > E_2$, and each component of the disjunction is handled by a clause. In practice, it is not necessary that all constraints in a constraint system can be negated, but only those that are required in a given program to answer some query.
- When the negation of a constraint cannot be expressed as a finite disjunction of constraints, we make a best effort to provide an *ad hoc* implementation. For example, for the equalities $v_i = t_i$ in the Herbrand domain CLP(H) that were added when normalizing clauses, we introduce a call $diff(v_i, t_i)$ to a disequality solver provided by the runtime environment. Negation of equality in CLP(H) can be expressed as a finite disjunction only for programs that can generate a finite number of ground terms.

Executable code for the dual program is generated by removing the external quantifiers (as in Horn clauses) and translating the universal quantifiers that were applied to local variables into a call to the predicate `forall(Var, Pred)`, provided by the s(CASP) runtime.

Example 1

Given the program

```

1  p(X):- q(X, Z), not r(X).           3  q(X, a):- X #> 5.
2  p(Z):- not q(X, Z), r(X).         4  r(X):- X #< 1.

```

its dual is shown below

```

1  % not p/1                               11 % not q/1
2  not p(A) :- not p_1(A), not p_2(A).    12 not q(A,B) :- not q_1(A,B).
3                                          13
4  not p_1(A) :- forall(B, not p_1(A,B)).  14 not q_1(A,B) :- B \= a.
5  not p_1(A,B) :- not q(A,B).           15 not q_1(A,B) :- A #=< 5.
6  not p_1(A,B) :- r(A).                 16
7                                          17 % not r/1
8  not p_2(A) :- forall(B, not p_2(A,B)).  18 not r(A) :- not r_1(A).
9  not p_2(A,B) :- q(B,A).               19 not r_1(A) :- A #>= 1.
10 not p_2(A,B) :- not r(B).

```

2.1.4 Sketch of the execution scheme

Queries to the original program extended with the dual rules are evaluated by a runtime environment. This is currently a metainterpreter written in Prolog that executes the algorithm described by Marple et al. (2012). This algorithm has similarities with SLD resolution, but it takes into account specific characteristics of ASP and the dual programs, such as the different kinds of loops, the denials, and the introduction of universal quantifiers in the body of the clauses. The main highlights of this algorithm are:

Loop handling: Two different cases are distinguished by Marple et al. (2012):

- When a call eventually invokes itself and there is an odd number of intervening negations (as in, e.g. $p :- q. \quad q :- \text{not } r. \quad r :- p.$), the evaluation fails (and backtracks) to avoid contradictions of the form $p \wedge \neg p$.
- When there is an even number of intervening negations, as in $p :- \text{not } q. \quad q :- r. \quad r :- \text{not } p.$ the metainterpreter generates several stable models, such as $\{p, \text{not } q, \text{not } r\}$ and $\{q, r, \text{not } p\}$.

Denials: The s(CASP) compiler automatically generates an auxiliary predicate that captures all the denials written by the programmer. This predicate is invoked during query evaluation to ensure that the returned models are consistent with the denials. The current implementation executes them at the end of the query evaluation, when a candidate model has been generated. It would however be possible to check them at appropriate points while the execution proceeds, in order to increase performance, as suggested by Marple and Gupta (2014).

The s(CASP) compiler also detects statically rules of the form $r :- q, \text{not } r.$ and introduces denials to ensure that the models satisfy $\neg q \vee r$, even if the atoms r or q are not needed to solve the query. This is done by building a dependency graph of the program and detecting the paths where this may happen, including across several calls. For the propositional case, such an analysis can be precise. For the non-propositional case, an over-approximation is calculated. In both cases, denials that are not used during program evaluation can be generated. These may impose a penalty in execution time, but are safe. Therefore, s(CASP) will state that the program

```
1 p :- not q.           2 q :- not p.           3 r :- not r.
```

has no stable models, regardless of the initial query.

Universal quantification: Universal quantifications in the body of the clauses are translated into the construction `forall(Var, Pred)`. This is implemented by the runtime environment by solving `Pred`, extracting the constraints attached to the quantified variables, and using these constraints negated to narrow the constraint store under which `Pred` is executed. This is iterated until failure or until the constraint store has an empty domain for the quantified variables. Arias et al. (2018) present this algorithm in more detail.

2.1.5 Execution with unsafe variables and uninterpreted Function Symbols

The code in Example 1 has variables that would be termed as *unsafe* in regular ASP systems: variables that appear in negated atoms in the body of a clause, but that do

not appear in any positive literal in the same body. Since s(CASP) synthesizes explicit constructive goals for these negated goals, the aforementioned code can be run as-is in s(CASP). The query `?- p(A).` generates three different models:

```
{ p(A | {A > 5}), q(A |{A > 5}, a), not r(A |{A > 5}) }
  A > 5

{ p(A | {A ≠ a}), not q(B |{B < 1}, A |{A ≠ a}), r(B |{B < 1}) }
  A ≠ a

{ p(a), not q(B | {B < 1}, a), r(B |{B < 1}) }
  A = a
```

where the notation $V \mid \{C\}$ for a variable V is intended to mean that V is subject to the constraints in $\{C\}$. The constraints $A = 5$, $A \neq a$, and $A = a$ correspond to the bindings of variable A that make the atom in the query `?- p(A)` belong to the stable model.

Another very relevant point where s(CASP) differs from ASP is in the possibility of using arbitrary uninterpreted function symbols to build, for example, data structures. While in mainstream ASP implementations these could give rise to an infinite grounded program, the s(CASP) execution model can deal with them similarly to Prolog, with the added power of the use of constructive negation in the execution and in the returned models.

Example 2

The predicate `member/2` below models the membership to a list as it is usual in (classical) logic programming. The query is intended to derive the conditions for one argument not to belong to a given list.

```
1 member(X, [X|Xs]).
2 member(X, [_|Xs]):- member(X, Xs).
3
4 list([1,2,3,4,5]).
5
6 ?- list(A), not member(B, A).
```

This program and query return in s(CASP) the following model and binding:

```
{ list([1,2,3,4,5]),
  not member(B | {B ≠ 1,B ≠ 2,B ≠ 3,B ≠ 4,B ≠ 5}, [1,2,3,4,5]),
  not member(B | {B ≠ 1,B ≠ 2,B ≠ 3,B ≠ 4,B ≠ 5}, [2,3,4,5]),
  not member(B | {B ≠ 1,B ≠ 2,B ≠ 3,B ≠ 4,B ≠ 5}, [3,4,5]),
  not member(B | {B ≠ 1,B ≠ 2,B ≠ 3,B ≠ 4,B ≠ 5}, [4,5]),
  not member(B | {B ≠ 1,B ≠ 2,B ≠ 3,B ≠ 4,B ≠ 5}, [5]),
  not member(B | {B ≠ 1,B ≠ 2,B ≠ 3,B ≠ 4,B ≠ 5}, []) }
  A = [1,2,3,4,5], B ≠ 1, B ≠ 2, B ≠ 3, B ≠ 4, B ≠ 5
```

That is, for variable B not to be a member of the list $[1,2,3,4,5]$ it has to be different from each of its elements.

In addition to default negation, s(CASP) supports classical negation to capture the explicit evidence that a literal is false, as mentioned in Section 2.1.1.

s(CASP) is implemented in Ciao Prolog (Hermenegildo *et al.* 2012) and is available at <https://gitlab.software.imdea.org/ciao-lang/scasp>.

2.1.6 *s(CASP) as a conservative extension of ASP*

The behavior of *s(CASP)* and ASP is the same for propositional programs. For programs featuring unsafe variables (legal in ASP, but not in mainstream ASP systems) or programs that could create data structures arbitrarily large or whose variable ranges are defined in infinite domains (either unbound or bound but dense), which are outside the standard domain of ASP systems as they cannot be finitely grounded, *s(CASP)* extends ASP in a consistent way. The domain of the variables is implicitly expanded to include a domain which can be potentially infinite.

Let us use an example introduced in Marple et al. (2017, Pag.12). We are interested in knowing whether $p(X)$ (for some X) is or not part of a stable model:

```
1 d(1).
2 p(X) :- not d(X).
```

The only constant in the program is 1, which is the only possible domain for X in the second clause. That clause is not legal for ASP, as X is an unsafe variable (Sect. 2.1.5). Adding a domain predicate call for it (i.e. adding $d(X)$ to the body of the second clause), makes its model be $\{d(1)\}$ (**not** $p(1)$ is implicit).

That second clause is however legal in *s(CASP)*. Making the query $?- p(X)$ returns the *partial* model $\{p(X|\{X \neq 1\}), \text{not } d(X|\{X \neq 1\})\}$ stating that $p(X)$ and **not** $d(X)$ are true when $X \neq 1$, which is consistent with, but more general than, the model given by ASP. As the model is partial, only the atoms (perhaps negated) involved in the proof for $?- p(X)$ appear in that model.

2.2 Circumscription

Circumscription (McCarthy 1980; Lifschitz 1985) is a technique to perform non-monotonic reasoning within the framework of first-order logic. Circumscription minimizes the extension⁴ of the predicates that we want to circumscribe. Intuitively, it aims at formalizing that the known objects in a certain class are *all* the objects that are in that class. EC theories require that some of their predicates are circumscribed to ensure that they can only be interpreted as they appear in the description of the scenario.

The following definition of circumscription is due to Lifschitz (1985):

Definition 1 (Circumscription)

Let $A(P, Z)$ denote a sentence, where P is a tuple of predicate constants and Z a tuple of function and/or predicate constants disjoint with P . The circumscription of P in $A(P, Z)$ with variable Z is defined as the second order sentence $A(P, Z) \wedge \neg \exists p, z (A(p, z) \wedge p < P)$. The $<$ symbol in the previous expression is defined as follows: if U and V are n -ary predicates, $U \leq V$ stands for $\forall x_1 \dots x_n (U(x_1, \dots, x_n) \rightarrow V(x_1, \dots, x_n))$. $U = V$ stands for $U \leq V$ and $V \leq U$ and $U < V$ stands for $U \leq V \wedge \neg (V \leq U)$. These definitions are extended to tuples of predicates in the obvious fashion.

$U \leq V$ expresses that the extension of U is a subset of the extension of V , and $U < V$ means that the extension of U is a proper subset of the extension of V . The circumscription of $A(P, Z)$ is a formula whose extension is the minimal extension of the predicates

⁴ The *extension* of a predicate is the set of tuples for which predicate is true.

in P that makes $A(P, Z)$ true, and in which the objects in Z are allowed to vary. It is expressed with $CIRC(A; P; Z)$ or $CIRC(A; P)$ if Z is empty.

Example 3

Let's take $A = P(a)$. Its minimization should express that P is only true for a :

$$CIRC(A; P) \equiv \forall x(P(x) \leftrightarrow x = a).$$

Example 4

Let's take $A = \neg P(a)$. We only have information about P not being true in a , which is the only constant that appears in A , but also *all* the constants that appear in the sentence. Its minimization expresses this as:

$$CIRC(A; P) \equiv \forall x \neg P(x).$$

Computing circumscriptions is in general a hard task. However, there are class of formulas (*separable formulas*) for which it has shown to be easy by Lifschitz (1985). Also, Ferraris *et al.* (2011) have shown that there is a close relationship between circumscription and stable models and indeed Lee and Palla (2012, Definition 2) proved that the stable model semantics coincides with the circumscription for a large class of formulas called *canonical formulas*. Non-canonical formulas can often be rewritten as canonical formulas, therefore expanding the range of coincidence of circumscription and stable model semantics.

2.3 Event calculus

EC (presented at length by, for example, Mueller (2014)) is a formalism for reasoning about events and change, of which there are several axiomatizations. There are three basic, mutually related, concepts in EC: *events*, *fluents*, and *time points*. An event is an action or incident that may occur in the world: for instance, a person dropping a glass is an event. A fluent is a time-varying property of the world, such as the altitude of a glass. A time point is an instant in time. Events may happen at a time point; fluents have a truth value at any time point or over an interval, and their truth values are subject to change upon the occurrence of an event. In addition, fluents may have (continuous) quantities associated with them when they are true.

For example, the status of a glass falling may be represented by two fluents: one that captures the fact that the glass is falling and another one that captures the height of the glass over the ground. The *event* of dropping a glass initiates the *fluent* that captures that the glass is falling and gives some initial value to its height. This height changes with time according to some formula. The event of catching the glass makes the fluent that reflects it is falling false and makes its height not to change, possibly until a further event takes place. An EC description consists of a domain narrative (Figure 1) and a universal theory (Figure 2). The domain narrative consists of the causal laws of the domain, the known events, and the fluent properties, and the universal theory is a conjunction of EC axioms that encode, for example, the inertia laws.

The original EC (OEC) was introduced by Kowalski and Sergot (1989). OEC has sorts for event occurrences, fluents, and time periods. In this paper we use the BEC formulated by Shanahan (1999) as presented by Mueller (2008a). BEC allows fluents to be released

Predicate	Meaning
$InitiallyN(f)$	fluent f is false at time 0
$InitiallyP(f)$	fluent f is true at time 0
$Happens(e, t)$	event e occurs at time t
$Initiates(e, f, t)$	if e happens at time t , f is true and not released from the commonsense law of inertia after t
$Terminates(e, f, t)$	if e occurs at time t , f is false and not released from the commonsense law of inertia after t
$Releases(e, f, t)$	if e occurs at time t , f is released from the commonsense law of inertia after t
$Trajectory(f_1, t_1, f_2, t_2)$	if f_1 is initiated by an event that occurs at t_1 , then f_2 is true at t_2
$StoppedIn(t_1, f, t_2)$	f is stopped between t_1 and t_2
$StartedIn(t_1, f, t_2)$	f is started between t_1 and t_2
$HoldsAt(f, t)$	fluent f is true at time t

Fig. 1. Basic event calculus (BEC) predicates
(e = event, f , f_1 , f_2 = fluents, t , t_1 , t_2 = timepoints).

- BEC1.** $StoppedIn(t_1, f, t_2) \equiv \exists e, t (Happens(e, t) \wedge t_1 < t < t_2 \wedge (Terminates(e, f, t) \vee Releases(e, f, t)))$
- BEC2.** $StartedIn(t_1, f, t_2) \equiv \exists e, t (Happens(e, t) \wedge t_1 < t < t_2 \wedge (Initiates(e, f, t) \vee Releases(e, f, t)))$
- BEC3.** $HoldsAt(f_2, t_2) \leftarrow Happens(e, t_1) \wedge Initiates(e, f_1, t_1) \wedge Trajectory(f_1, t_1, f_2, t_2) \wedge \neg StoppedIn(t_1, f_1, t_2)$
- BEC4.** $HoldsAt(f, t) \leftarrow InitiallyP(f) \wedge \neg StoppedIn(0, f, t)$
- BEC5.** $\neg HoldsAt(f, t) \leftarrow InitiallyN(f) \wedge \neg StartedIn(0, f, t)$
- BEC6.** $HoldsAt(f, t_2) \leftarrow Happens(e, t_1) \wedge Initiates(e, f, t_1) \wedge t_1 < t_2 \wedge \neg StoppedIn(t_1, f, t_2)$
- BEC7.** $\neg HoldsAt(f, t_2) \leftarrow Happens(e, t_1) \wedge Terminates(e, f, t_1) \wedge t_1 < t_2 \wedge \neg StartedIn(t_1, f, t_2)$

Fig. 2. Formalization of BEC theory as presented by Mueller (2014).

from the commonsense law of inertia via the *Release* predicate, and adds the ability to represent continuous change via the *Trajectory* predicate. Figure 2 summarizes the seven axioms of the BEC theory. An explanation of these axioms follows:

- **Axiom BEC1.** A fluent f is stopped between time points t_1 and t_2 iff it is terminated or released by some event e that occurs after t_1 and before t_2 .
- **Axiom BEC2.** A fluent f is started between time points t_1 and t_2 iff it is initiated or released by some event e that occurs after t_1 and before t_2 .
- **Axiom BEC3.** A fluent f_2 is true at time t_2 if a fluent f_1 initiated at t_1 does not finish before t_2 and it makes fluent f_2 be true.⁵
- **Axiom BEC4.** A fluent f is true at time t if it is true at time 0 and is not stopped on or before t .

⁵ For implementation convenience, and without loss of expressiveness, we assume that argument t_2 in $Trajectory(f_1, t_1, f_2, t_2)$ is not a time difference w.r.t. t_1 , but an absolute time after t_1 .

- **Axiom BEC5.** A fluent f is false at time t if it is false at time 0 and it is not started on or before t .
- **Axiom BEC6.** A fluent f is true at time t_2 if it is initiated at some earlier time t_1 and it is not stopped before t_2 .
- **Axiom BEC7.** A fluent f is false at time t_2 if it is terminated at some earlier time t_1 and it is not started on or before t_2 .

3 From event calculus to *s(CASP)*

3.1 Circumscription in *s(CASP)*

Circumscription is applied to EC domain narratives, and as a result, the events that happen and their effects are only those explicitly defined. As mentioned before, the definition of a given scenario (its *narrative* part) states the basic actions and effects using the predicates in Figure 1. Let us consider example 14 by Mueller (2014), which reasons about the turning on and off of a light switch:

$$\begin{aligned} \text{Happens}(e, t) \equiv & (e = \text{TurnOn} \wedge t = 2) \vee \\ & (e = \text{TurnOff} \wedge t = 4). \end{aligned}$$

Assuming circumscription, we can write the axioms $\text{Happens}(\text{TurnOn}, 2)$ and $\text{Happens}(\text{TurnOff}, 4)$, instead of writing the previous formula, while ensuring that there are no events (resp., effects) other than those stated. That is, we can prove that the light is off at $t = 6$, because we can prove the absence of an event turning the light on between $t = 4$ and $t = 6$. If we need to add a new event (e.g. $\text{Happens}(\text{TurnOn}, 5)$) we only have to add it, instead of modifying the formula that expresses the circumscription of the narrative. It has to be noted that in this case the circumscribed formula is simple, but the circumscription of more complex formulas is more involved.

Similarly, since EC assumes the circumscription of the rest of the predicates defined in the narrative (*InitiallyN*, *InitiallyP*, *Initiates*, *Terminates*, *Releases*, and *Trajectory*), we have that the explicitly known effects of events are the only effects of events.

EC does not assume circumscription for the underlying theory. Since BEC1 and BEC2 are definitions which can be expanded wherever they appear, they do not need to be taken care of specially.

However, for the rest of the BEC theory axioms in Figure 2, from BEC3 to BEC7, we **cannot** apply this assumption because they are not circumscribed and they are implications: *HoldsAt* is true *if* the body of one of its corresponding axioms holds. This means that, if we cannot infer that $\text{HoldsAt}(f, t)$ is true, then, we **cannot** deduce that $\neg\text{HoldsAt}(f, t)$ is true. Instead, the falsehood of $\text{HoldsAt}(t, f)$ must be inferred through the axioms BEC5 and BEC7 if this is supported by direct evidences defined in the narrative. It is important to note that if the narrative describes a scenario in which it is possible to deduce that a fluent f is true and false at the same point in time t , this narrative is inconsistent and there are no valid models for this scenario. Furthermore, if a given narrative is described in terms of the truth or falsehood of a fluent f at some point in time t , but we are not able to decide this value, then there are multiple valid models because we have that $\text{HoldsAt}(t, f) \vee \neg\text{HoldsAt}(t, f)$.

Let us describe below how we use $s(\text{CASP})$ to compute the circumscription based on predicate completion (Mueller 2014; Lee and Palla 2012) while we are able to express the truth/falsehood of *HoldsAt* using classical negation.

3.2 Modeling BEC with $s(\text{CASP})$

Two key factors contribute to $s(\text{CASP})$'s ability to model EC: the preservation of non-ground variables during the execution and the integration with constraint solvers.

Treatment of variables in $s(\text{CASP})$: Thanks to the use of non-ground variables, $s(\text{CASP})$ is able to directly model EC axioms that would otherwise require “unsafe” rules (Section 2.1.5). Let us take, for example, rule BEC4 of Figure 2. In a straightforward encoding (see Figure 3), the parameter t (T in the code), which appears in the head and does not appear in a positive literal in the body (i.e. it only appears in $\neg \text{StoppedIn}(0, f, t)$) would be classified as unsafe by a mainstream ASP system.

It may be argued that a way to overcome this issue would be to translate the negation in the axioms as classical negation, that is, $\text{-stoppedIn}(0, F, T)$. However, this would need to generate a negation of the two clauses of $\text{stoppedIn}/3$ which would in turn need the (classical) negations of $\text{terminates}/3$, $\text{happens}/2$, and $\text{releases}/3$. However, these narrative predicates do not come with a definition stating when they are false, in the classical sense of negation.

```

1  %% BEC1                                29  %% BEC4
2  stoppedIn(T1,F,T2) :-                 30  holdsAt(F,T) :-
3      T1 #< T, T #< T2,                  31      0 #< T,
4      terminates(E,F,T),                 32      initiallyP(F),
5      happens(E,T).                       33      not stoppedIn(0,F,T).
6
7  stoppedIn(T1,F,T2) :-                 34
8      T1 #< T, T #< T2,                  35  %% BEC5
9      releases(E,F,T),                   36  -holdsAt(F,T) :-
10     happens(E,T).                       37      0 #< T,
11
12  %% BEC2                                38      initiallyN(F),
13  startedIn(T1,F,T2) :-                 39      not startedIn(0,F,T).
14     T1 #< T, T #< T2,                   40
15     initiates(E,F,T),                   41  %% BEC6
16     happens(E,T).                       42  holdsAt(F,T2) :-
17
18  startedIn(T1,F,T2) :-                 43     T1 #< T2,
19     T1 #< T, T #< T2,                   44     initiates(E,F,T1),
20     releases(E,F,T),                   45     happens(E,T1),
21     happens(E,T).                       46     not stoppedIn(T1,F,T2).
22
23  %% BEC3                                47
24  holdsAt(F2,T2) :-                     48  %% BEC7
25     initiates(E,F1,T1),                 49  -holdsAt(F,T2) :-
26     happens(E,T1),                       50     T1 #< T2,
27     trajectory(F1,T1,F2,T2),            51     terminates(E,F,T1),
28     not stoppedIn(T1,F1,T2).            52     happens(E,T1),
                                           53     not startedIn(T1,F,T2).
                                           54
                                           55  %% Consistency (automatically added)
                                           56  :- -holdsAt(F,T), holdsAt(F,T).

```

Fig. 3. Basic Event Calculus (BEC) modeled in $s(\text{CASP})$.

An ASP solver such as *clingo* (Gebser *et al.* 2014) will not be able to directly process unsafe rules like this. The standard approach to fix unsafe rules is to add a positive literal defining the domain of the unsafe variable (T in this case), but this is not feasible if we want to maintain the property that T represents time and is therefore not discrete and not finite. On the other hand, the top-down execution strategy of s(CASP) makes it possible to keep logical variables both during execution and in answer sets and therefore free (logical) variables can be handled in heads and in negated atoms.

Integration with constraint solvers: The s(CASP) system has a generic interface to enable plugging in constraint solvers. s(CASP) currently includes the CLP(\mathbb{Q}) linear constraints solver by Holzbaur (1995), that supports the arithmetic constraints $<$, $>$, $=$, \leq , \geq . This is used to implement the definitions and axioms of BEC that require comparisons of points in (continuous) time and to solve the equations that arise from these comparisons. The selection of CLP(\mathbb{Q}) instead of the faster CLP(\mathbb{R}) is motivated by soundness reasons. Since CLP(\mathbb{R}) uses internally floating-point numbers, rounding and approximations compromise accuracy and termination of some code. On the other hand, CLP(\mathbb{Q}) represents rational numbers exactly and therefore it should not introduce any calculation error. One example in which the use of floating-point numbers would be inadequate is the code in Figure 5, which uses the factor $\frac{4}{3}$ and that does not have an exact floating-point representation.

The s(CASP) infrastructure is however parametric w.r.t. the underlying constraint solver, and other implementations of constraint domains can be implemented and plugged in if necessary. As an example, s(CASP) includes a solver for disequality in the Herbrand domain (Sec. 2.1.3), which is necessary to generate the dual of almost any interesting program.

3.3 Translating the BEC axioms into s(CASP)

Our translation of the BEC axioms into s(CASP) is related to the translation used by the systems EC2ASP and F2LP (Lee and Palla 2012; 2020). We differ in three aspects that improve performance for a top-down system, fully use s(CASP)'s ability to treat unbound variables, and do not compromise the soundness of the translation (according to the proofs presented by Lee and Palla (2012)). These are: *the treatment of rules with negated heads*, *the possibility of generating unsafe rules*, and *the use of constraints over dense domains* (rationals, in our case). We describe below, with the help of a running example, the translation that turns logic statements (as found in BEC) into an s(CASP) program. The code corresponding to the translations of the axioms of BEC in Figure 2 can be found in Figure 3. s(CASP) code follows the syntactical conventions of logic programming: constants (including function names) and predicate symbols start with a lowercase letter and variables start with an uppercase letter. In addition, numerical constraints are written as constraints in s(CASP), (e.g. $\#<$) to make it clear that they do not correspond to Prolog's arithmetic comparisons:

- **Atoms and Constants:** Their names are preserved. *Uniqueness of Names* (Shanahan 1999) is assumed by default (and enforced) in logic programming.
- **Constraints:** Predicates that represent constraints (e.g. on time) are directly translated to their counterparts in s(CASP). For example, $t_1 < t_2$ becomes

T1 #< T2, which is handled by the CLP(Q) solver. The translation is parameterized on the constraint domain.

- **Definitions:** The axiomatization of BEC uses definitions of the form $D(x) \equiv \exists y B(x, y)$, where $B(x, y)$ is a conjunction of (possibly negated) atoms, disjunctions of atoms, and constraints (e.g. BEC1). The use of these definitions makes it easier to build and reuse conceptual blocks out of basic predicates. They are however not strictly necessary *per se* as predicates. For performance reasons we treat them as if they were written as $\forall x(D(x) \leftarrow \exists y B(x, y))$, following the work of Lee and Palla (2020). Since D , as mentioned above, is given a name as a convenience to write the BEC axioms, we can ignore its truth value in the (partial) models that s(CASP) generates because if it were expanded where it is used, it would have disappeared. Therefore, the models returned using implication or equivalence for the definition of D are the same.
- **Rules with Positive Heads:** A rule (e.g. BEC6)

$$\forall x(H(x) \leftarrow \exists y(A(y) \wedge \neg B(x, y) \wedge x < y)),$$

where $x < y$ is a constraint, is translated into

$$1 \quad \mathbf{h(X)} \text{ :- } \mathbf{X \#< Y}, \mathbf{a(Y)}, \text{ not } \mathbf{b(X, Y)}.$$

The constraint $\mathbf{X \#< Y}$ could be placed anywhere in the clause. However, in top-down evaluation schemes, the general recommendation is to execute constraints as soon as possible to use a *constraint and generate* mechanism instead of *generate and test* in order to improve performance. In the very common case where user-level constraint operations are translated into constraint propagation, which is usually required to be deterministic, executing them earlier in the tree simplifies the constraints without performing internal search or creating search branches. On the other hand, by constraining the domains of the variables earlier, the size of the search trees needed by calls to user predicates ($\mathbf{a(Y)}$, $\text{not } \mathbf{b(X, Y)}$ in this case) are reduced.

- **Rules with Negated Heads:** BEC rules 5 and 7 infer negated heads $\neg \text{HoldsAt}(f, t)$ while rules 4 and 6 infer positive heads $\text{HoldsAt}(f, t)$, that is, they follow, respectively, the scheme

$$\forall x(\neg H(x) \leftarrow \exists y A(x, y)) \quad \wedge \quad \forall x(H(x) \leftarrow \exists y B(x, y)).$$

The standard approach to translate rules with negated heads is to convert them into denials (Lee and Palla 2012) such as $\text{:- } \mathbf{a(X, Y)}, \mathbf{h(X)}$. Our approach is to create instead the atom $\neg \mathbf{h(X)}$ to denote (Section 2.1.1) the negation of $\mathbf{h(X)}$ and a rule that captures the explicit evidence that $\mathbf{h(X)}$ is false:

$$1 \quad \neg \mathbf{h(X)} \text{ :- } \mathbf{a(X, Y)}.$$

This makes it possible to invoke $\neg \mathbf{h(X)}$ as a regular predicate in a top-down execution. The compiler will additionally (Section 2.1.1) generate denials to ensure that $\neg \mathbf{h(X)}$ and $\mathbf{h(X)}$ cannot be simultaneously true. Therefore, s(CASP) will detect an inconsistency if both $\text{HoldsAt}(f, t)$ and $\neg \text{HoldsAt}(f, t)$ can be simultaneously derived from the narrative. Having rules stating explicitly when $\neg \text{HoldsAt}(f, t)$ can be derived makes it possible to query for it, as some BEC rules need. We will later see how this is connected with the translation of the narrative.

- **Rules with Disjunctive Bodies:** A rule (e.g. BEC1)

$$\forall x[H(x) \leftarrow \exists y((A(x, y) \vee B(x, y)) \wedge C(x, y))]$$

is translated into two separate clauses:

- 1 $h(X) \text{ :- } a(X, Y), c(X, Y).$
- 2 $h(X) \text{ :- } b(X, Y), c(X, Y).$

3.4 Translation of the narrative

Every basic BEC predicate $P(x)$ from the narrative (Figure 1) is translated into s(CASP) as a set of rules of the form

$$P(x) \leftarrow \gamma,$$

where the body γ is a conjunction of atoms, negated atoms, and constraints that can be trivially *true* in some cases. The rules corresponding to the same $P(x)$ are assumed to state **all** the cases where $P(x)$ is true.

Throughout this section we will use the light scenario problem from Mueller (2014, example 14): a two-color bulb lamp can be switched on and off. We want to be able to answer when the light is on (and its color in this case) and off. We will present the BEC narrative and its translation into s(CASP), available in full in Figure 4.

Events: Let us consider the description below:

$$\begin{aligned} Happens(e, t) \equiv & (e = TurnOn \wedge (t = 2 \vee t = 5)) \vee \\ & (e = TurnOff \wedge t = 4). \end{aligned}$$

It states that the *TurnOn* event will happen at time $t = 2$ and $t = 5$, and that *TurnOff* will happen at $t = 4$. As we mentioned before, since EC assumes circumscription, it is equivalent to the axioms $Happens(TurnOn, 2)$, $Happens(TurnOff, 4)$ and $Happens(TurnOn, 5)$, which are translated as facts (lines 1-3 of Figure 4).

Event Effects: The effects of events are represented using the predicates $Initiates(e, f, t)$ and $Terminates(e, f, t)$. In our example, when the event *TurnOn* happens, the light is put in *on* status; similarly, when the event *TurnOff* happens, the light will not be *on*, *red* or *green*:

$$\begin{aligned} Initiates(e, f, t) & \equiv (e = TurnOn \wedge f = On) \\ Terminates(e, f, t) & \equiv (e = TurnOff \wedge (f = On \vee f = Red \vee f = Green)). \end{aligned}$$

```

1 happens(turn_on, 2).
2 happens(turn_off, 4).
3 happens(turn_on, 5).
4
5 initiates(turn_on, on, T).
6
7 terminates(turn_off, on, T).
8 terminates(turn_off, red, T).
9 terminates(turn_off, green, T).
10 trajectory(on, T1, red, T2) :-
11     T1 #< T2,
12     T2 #< T1 + 1.
13 trajectory(on, T1, green, T2) :-
14     T2 #>= T1 + 1.
15
16 releases(turn_on, red, T).
17 releases(turn_on, green, T).
```

Fig. 4. Narrative of the light scenario modeled in s(CASP).

In both cases, this can happen at any time t , and the translation again becomes facts (lines 5-9 in Figure 4)

Release from Inertia: When turned on, the light emits red light for one second and after that it starts to emit green light. *Trajectory* expresses how this change depends on the time elapsed since an event occurrence.

$$\begin{aligned} \text{Trajectory}(s, t_1, c, t_2) \quad \equiv \quad & (s = \text{On} \wedge c = \text{Red} \wedge t_1 < t_2 \wedge t_2 < t_1 + 1) \vee \\ & (s = \text{On} \wedge c = \text{Green} \wedge t_2 \geq t_1 + 1). \end{aligned}$$

The translation is in lines 10-14 of Figure 4. *Releases* states that the color of the light is released from the commonsense law of inertia. After a fluent is released, its truth value is not determined by BEC and it can change. Thus, there may be models in which the fluent is true, and models in which the fluent is false:

$$\text{Releases}(e, f, t) \quad \equiv \quad (e = \text{TurnOn} \wedge (f = \text{Red} \vee f = \text{Green})).$$

Note that the time parameter t appears only in the head (it is universally quantified). Releasing a fluent (lines 16 and 17 in Figure 4) frees it up so that other axioms in the domain description can be used to determine its truth value, thus allowing us to represent continuous change of the fluent.

State Constraints: State constraints usually contain $\text{HoldsAt}(f, t)$ or $\neg\text{HoldsAt}(f, t)$ and represent restrictions on the models. In our running example, a light cannot be red and green at the same time:

$$\neg\text{HoldsAt}(\text{Red}, t) \vee \neg\text{HoldsAt}(\text{Green}, t).$$

Note that this is logically equivalent to $\forall t \neg(\text{HoldsAt}(\text{Red}, t) \wedge \text{HoldsAt}(\text{Green}, t))$ and is translated as `:- holdsAt(red, T), holdsAt(green, T)`. Adding this denial to the program in Figure 4 does not change its models. However, if we change the trajectory definition for the red light stating $t_2 \leq t_1 + 1$ instead of $t_2 < t_1 + 1$ the state constraint is violated at $t_2 = t_1 + 1$ and therefore, there are no valid models.

A Note on Using $\neg\text{HoldsAt}(f, t)$ in the Narrative: The narrative predicates may depend on what the BEC theory can deduce. In other words, $\text{HoldsAt}(f, t)$ or $\neg\text{HoldsAt}(f, t)$ can (perhaps indirectly) be used in the body γ of a narrative rule (see an example in Figure 5). $\text{HoldsAt}(f, t)$ can be invoked directly, but $\neg\text{HoldsAt}(f, t)$ ought to be called using classical negation, For example, `-holdsAt(F, T)`. As we mentioned before, the reason is that circumscription is not applied to the EC axioms and we can deduce only the truth (or falsehood) of HoldsAt when we have direct evidence of either of them — that is, what the positive (`holdsAt(F, T)`) and negative (`-holdsAt(F, T)`) heads provide.

As we mentioned before, the consistency rule (line 56 in Figure 3) introduced by the compiler of s(CASP) would ensure that `holdsAt(F, T)` and `-holdsAt(F, T)` are mutually exclusive by flagging an inconsistency. Note that if `not holdsAt(F, T)` succeeds, then `holdsAt(F, T)` is false but it does not imply that `-holdsAt(F, T)` is true (i.e. `-holdsAt(F, T) ⇒ not holdsAt(F, T)`), but it is **not** the case that `not holdsAt(F, T) ⇒ -holdsAt(F, T)`. Symmetrically for `not -holdsAt(F, T)` we have that `holdsAt(F, T) ⇒ not -holdsAt(F, T)`.

```

1  #include bec_theory.
2
3  max_level(10) :- not max_level(16).
4  max_level(16) :- not max_level(10).
5
6  initiallyP(level(0)).
7
8  happens(overflow,T).
9  happens(tapOn,5).
10
11 initiates(tapOn,filling,T).
12 terminates(tapOff,filling,T).
13 initiates(overflow,spilling,T):-
14     max_level(Max),
15     holdsAt(level(Max), T).
16
17 releases(tapOn,level(0),T):-
18     happens(tapOn,T).
19 trajectory(filling,T1,level(X2),T2):-
20     T1 #< T2,
21     X2 #= X + 4/3*(T2-T1),
22     max_level(Max),
23     X2 #=< Max,
24     holdsAt(level(X),T1).
25 trajectory(filling,T1,overlimit,T2):-
26     T1 #< T2,
27     X2 #= X + 4/3*(T2-T1),
28     max_level(Max),
29     X2 #> Max,
30     holdsAt(level(X),T1).
31
32 trajectory(spilling,T1,leak(X),T2):-
33     holdsAt(filling, T2),
34     T1 #< T2,
35     X #= 4/3*(T2-T1).

```

Fig. 5. Encoding of an Event Calculus narrative with continuous change.

In previous implementations of EC, such as F2LP, reasoning about the falsehood of $HoldsAt(f,t)$ can be made using only the default negation, implemented as negation as failure (i.e. `not/1`). Therefore, the presence of classical negation in s(CASP) not only increases the expressive power from the point of view of the programmer (by, e.g. determining whether a fluent is or not active at some point in time), but it also ensures correctness in those cases where $\neg holdsAt(F,T) \neq not\ holdsAt(F,T)$, as inconsistent models are discarded.

3.5 Continuous change: A complete encoding

We consider now an extension of the water tap example by Shanahan (1999), where we define two possible worlds and added a triggered fluent to describe the ability of s(CASP) to model complex scenarios. In this example, a water tap fills a vessel, whose capacity is either 10 or 16 units, and when the level of water reaches the bucket rim, it starts spilling. Let us present the main ideas behind its encoding, available in Figure 5.

Continuous Change: The fluent $Level(x)$ represents that the water is at level x in the vessel. The first *Trajectory* formula (lines 19-24) determines the time-dependent value of the $Level(x)$ fluent, which is active as long as the *Filling* fluent is true and the rim of the vessel is not reached. The second *Trajectory* formula (lines 25-30) allows us to capture the fact that the water reached the rim of the vessel and overflowed. Note the $\frac{3}{4}$ factor that relates time and water amount. As mentioned before, if the underlying solver approximates the numerical operations, water levels could be miscalculated and therefore the answers to some queries could be wrong. The use of $CLP(\mathbb{Q})$ prevents this.

Uniqueness of $Level(x)$ A relevant question is whether the fluent $Level(x)$ could take two different values at the same time. Intuitively, it should not, because if we are modeling faithfully a physical system evolving under a series of events, the level should be unique at any point in time. Note, however, that if this were to happen, it would be because the

narrative does not correspond to the reality. The specification given by Shanahan (1999) includes explicitly an axiom

$$\text{HoldsAt}(\text{Level}(x_1), t) \wedge \text{HoldsAt}(\text{Level}(x_2), t) \rightarrow x_1 = x_2,$$

to avoid this situation.

We did not include it for two reasons: on the one hand, a careful inspection of the narrative reveals that this could not be the case. On the other hand, if this axiom were necessary, the model without it would allow the simultaneous existence of two alternative water levels. That raises the question whether the model is correct, as the narrative would in that case allow two different, diverging events happen at the same time, or the trajectories would allow two inconsistent fluents overlap. Just stating this inconsistency is of little value in practice, as it does not help catch errors in the model and it removes possible correct states.⁶

Triggered Fluent: The fluent *Spilling* is triggered (lines 13-15) when the water level reaches the rim of the vessel. As a consequence, the *Trajectory* formula in lines 32-35 starts the fluent *Leak*(x) and captures the amount of water leaked while the fluent *Spilling* holds.

Alternative Worlds: As we mentioned in Section 2.1.4, the presence of even loops generates different worlds. In our implementation, the clauses in lines 3-4 force the vessel capacity to be either 10 or 16, and therefore, they create two possible worlds/models: $\{\text{max_level}(10), \text{not max_level}(16), \dots\}$ and $\{\text{max_level}(16), \text{not max_level}(10), \dots\}$.

Different worlds can be used to model alternative scenarios where an event may happen in one world and not in another. For this, a keyword `#abducible` is provided as a shortcut in s(CASP). We will use it in Sec. 4.2, below.

4 Examples and evaluation

The benchmarks used in this section are available at <http://www.cliplab.org/papers/EC-sCASP-TPLP2020/>. They were run on a macOS 10.15.7 laptop with an Intel Core i5 at 2GHz.

4.1 Deduction

Deduction determines whether a state of the world is possible given a theory and an initial narrative. We can perform deduction in BEC for the previous examples through queries to the corresponding s(CASP) program. For the *lights* scenario (Figure 4):

`?- holdsAt(on,3)` succeeds: it states that the light is on at time 3.

⁶ Let us note that EC does not preclude a fluent to have different associated values simultaneously: it is only the semantics of this fluent that may disallow it. As an example, let us take the fluent *Occupied*(n) that expresses that the seat n in a theater is occupied. Obviously, several instances for different n can be true at the same time.

```
?- -holdsAt(on,4.5) succeeds: the light is not on at time 4.5.7
?- holdsAt(F,3) is true in one stable model containing holdsAt(green,3) and
   holdsAt(on,3), meaning that the light is on and green at time 3.
```

Additionally, as we mentioned in Section 3.4, using the default negation `not/1` we can check the absence of a proof for `holdsAt/2` and `-holdsAt/2`. However, there are time points, for example, at time 1, where neither the truth nor the falsehood for the fluent representing that the light is on can be deduced from the program. Therefore, the queries `?- not holdsAt(on,1)` and `?- not -holdsAt(on,1)` would both succeed.

Finally, let us use the water level scenario (Figure 5) to make queries involving time and the water level (that are continuous physical quantities):

```
?- holdsAt(level(H),15/2) is true when H=10/3.
?- holdsAt(level(10/3),T) is true when T=15/2.
```

Note that, as explained with more detail in the *Evaluation* subsection below, *s(CASP)* can operate and answer correctly queries involving rationals (and, in general, dense domains as long as they are supported by the underlying constraint solver) without having to modify the original program to introduce domains for the relevant variables or to *scale* the constants to convert rationals into integers.

4.2 Abduction

Abductive reasoning can be used to determine a sequence of events/actions that reaches a given state. In the case of ASP, actions are naturally captured as the set of atoms that are true in a model that includes the initial and final states, and that are consistent with the BEC theory. For the water scenario (Figure 5), let us assume we want to determine whether the water can reach a level of 12 at time 14. The query `?- holdsAt(level(12),14)` will return a single model with a vessel size of 16 and the rest of the atoms in the model capturing what must (not) happen to reach this state. For EC, the relevant atoms are those related to the events that happen (as these trigger fluents according to the model rules in `initiates/3`, `releases/3`, and `trajectory/4`) and, to keep track of the state of the system, the atom `holdsAt/2` and its classical negation. If, for the query mentioned above, we restrict the model to these atoms,⁸ we obtain the following filtered model:

```
{ initiallyP(level(0)), not happens(tapOn,D | {D #> 0,D #< 5}),
  holdsAt(level(0),5), happens(tapOn,5),
  not happens(tapOff,F | {F #> 5,F #< 14}), holdsAt(level(12),14) }
```

where we have listed the atoms in increasing order of time stamp.

This subset of the model states what must happen to reach a state with a water level of 12 at time $t = 14$. It captures events that must be observed (e.g. `happens(tapOn,5)`, meaning that the water tap has to open at time $t = 5$), what is the observable state of the fluents (`holdsAt(level(12),14)`, meaning that the water level at $t = 14$ is 12), but also

⁷ The decimal number 4.5 is automatically converted to rational representation.

⁸ This can be done automatically by placing `#show` directives in the source code.

what events **must not** happen (`not happens(tapOn, D | {D #> 0, D #< 5})`), meaning that the tap must not be opened between $t = 0$ and $t = 5$, or `not happens(tapOff, F | {F #> 5, F #< 14})`, meaning that the water tap must not be closed between $t = 5$ and $t = 14$).

Additionally, since `s(CASP)` only generates partial models, it does not contain atoms that express actions that are not necessary for the conclusion, that is, the plan does not contain references to actions (either positive or negative) that do not interfere with the final state. For example, it does not state that the event of closing the tap must or must not happen between $t = 0$ and $t = 5$: whether this event happens or not is immaterial for the final result. Other abductive tasks can be performed: adding the directive `#abducible` to the fact `happens(tapOn, 5)`, we specify that it is possible (but not necessary) for the tap to be open at time 5. As we mentioned in Section 3.5, this directive is translated into code that creates different worlds/models. For the previous query, `?- holdsAt(level(L), 14)` (that determines the level of water at $t = 14$), we obtain two alternative partial models:

- One containing the literal `happens(tapOn, 5)` meaning that the tap is open at $t = 5$, and therefore, the resulting model is the previous one.
- Another one containing `{ holdsAt(level(0), 14), initiallyP(level(0)), not happens(tapOn, G | {G #> 0, G #< 5}), not happens(tapOn, 5), not happens(tapOn, E | {E #> 5, E #< 14}) }` meaning that the tap is **not** open at $t = 5$ (and neither for $0 < t < 5$, nor for $5 < t < 14$), and therefore, the water level at $t = 14$ remains equal to 0, which causes the literal `holdsAt(level(0), 14)` be part of the model.

Note that `s(CASP)` determined the truth value of *Happens* and, more importantly, performed constraint solving to infer the time ranges during which some events ought (and ought not) to take place, represented by the negated atoms in the models inferred by constructive negation. Since all relevant atoms have a time parameter, they actually represent a *timed plan*. Due to the expressiveness of constraints, this plan contains information on time points when events must (not) happen and also on time *windows* (sometimes in relation with other events) during which events must (not) take place. Note that it would be impossible to (finitely) represent this interval with ground atoms, as it corresponds to an infinite number of points.

4.3 Evaluation

A direct performance comparison of our implementation of BEC in `s(CASP)` with implementations in other systems may not be meaningful: most previous systems implement *discrete* Event Calculus (DEC) and they do not support continuous quantities. Since offering this support is one of the key points of our proposal, giving up on it and comparing with an implementation DEC in `s(CASP)` is pointless and defeats the main purpose of this piece of work. We will then have to compare BEC in `s(CASP)` with how a system that implements DEC can be used to approximate the results we can give.

One of the ASP-based tools that support DEC is F2LP, an ASP-based system that according to Lee and Palla (2012) “outperforms DEC reasoner (Mueller 2008b)”, reported there as the more efficient SAT-based implementation. F2LP is a tool that executes DEC

by turning first-order formulas under the stable model semantics into a logic program without constraints that is then evaluated using an ASP solver.

Our first evaluation compares the light scenario in Figure 4 running under *s(CASP)* with the F2LP translation under *clingo* 5.1.1, the current version of the state-of-the-art ASP system. Since the directive `#domain` is no longer available in *clingo*, we adapted the translation of F2LP adding `timestep(1..10)` and `timestep/1` to make the clauses safe (Appendix A of the supplementary material accompanying the paper at the TPLP archive). While under *s(CASP)* we can reason about time points in an unbounded, dense domain, the encoding used by F2LP makes time belong to the integers (in particular, to the interval from 1 to 10, with the previous `timestep/1` fact). Therefore, since the light can be red for $t > 2, t < 3$ and for $t > 5, t < 6$,⁹ there are no **integer** time points between 1 and 10 where the emitted light is red. Therefore, the query `?- holdsAt(red,T)` does not return any model under *clingo*, while the execution of the same query under *s(CASP)* returns the constraint conjunctions `T#>2, T#<3` and `T#>5, T#<6`.

In order to determine at what time point the red light is on using *clingo*, we modified the program generated by F2LP to refine the `timestep` domain with `timestep(1..10*P):- precision(P)`, where the new predicate `precision(P)` makes it possible to have a finer grain for the possible values of `timestep` by increasing the value of `P`. With this modification, it is possible to check if the light is red at time $t = 5.9$ with *clingo* by stating that we want a precision of tenths (using `precision(10)`) and modifying the queries accordingly, for example, using `?- holdsAt(red,59)`. Similarly, using `p=100` in `precision(P)` it is possible to check for $t = 5.99$ by querying `?- holdsAt(red,599)`, and so on.

This change (see Appendix B of the supplementary material accompanying the paper at the TPLP archive, for the complete program) is in principle not difficult to perform, but it undoubtedly obfuscates the resulting program (and for more complex narratives it would be harder, impractical, or even infeasible), and also impacts negatively its performance. Table 1a shows how the additional precision necessary in the F2LP encoding increases the execution run-time of *clingo* by orders of magnitude. On the other hand, *s(CASP)* does not have to adapt its encoding/queries and its performance does not change.

The second benchmark is the water level scenario in Figure 5. The physical model needs continuous quantities but, as it happened in the previous case, its execution with *clingo* forces the level of water to be expressed using integers. For example, the query `?- holdsAt(level(11),T)` returns the value `T=53/4` ($= 13.25$) using *s(CASP)*,¹⁰ while the execution under *clingo* fails: the level of water is $h = 10.\overline{6}$ at time $t = 13$ and $h = 12$ at time $t = 14$; therefore there is no integer point in time t with a level of water $h = 11$.

As before, we modified the program generated by F2LP by adding the predicate `precision(P)` to specify the precision of the program by scaling the numbers *clingo* deals with, which is necessary to determine at what time the level of the water is 11. For this example, the resulting encoding (see Appendix C of the supplementary material accompanying the paper at the TPLP archive) is more complex than for the previous

⁹ When turned on (at $t = 2$ and $t = 5$), the light emits red light for one second (see Section 3.4).

¹⁰ *s(CASP)* returns a rational number because it uses the rational constraint solver `CLP(Q)` in order not to lose precision, as would happen if using floating-point numbers. *s(CASP)* can however output them in decimal notation by using the command-line flag `-r`.

Table 1. *Comparative table of s(CASP) and F2LP+clingo*

(a) Run time (ms) comparison for the light scenario

Queries	s(CASP)	F2LP+clingo and precision	
?- holdsAt(red,5.9)	228	82	10
?- holdsAt(red,5.99)	240	8,364	100
?- holdsAt(red,5.999)	226	> 5 min.	1000

(b) Run time (ms) comparison for the water tap scenario

Queries	s(CASP)	F2LP+clingo and precision	
?- holdsAt(level(11),T)	301	475	10*
		77,305	100

(*) With this precision the value returned by *clingo* is wrong.

benchmark. Additionally, this example shows that the scaling value depends on the particular benchmark: `precision(10)` is not fine-grained enough to capture the solution, and we have to go up to `precision(100)` to obtain a model with `T=1325` (corresponding to $t = 13.25$, the correct value) for the query `holdsAt(level(1100),T)`.

An undesirable effect of rounding in ASP is that rounding may not only make programs fail, but it may make them succeed with wrong answers. For example, in the water level example, with `precision(10)`, the query `holdsAt(level(110),T)` holds with `T=133` (which would correspond to $t = 13.3$). This value is not right, and it is due to arithmetic rounding performed during the program execution.

Also, note that increasing P as powers of 10, as we have done, may not always work: scaling units to make an answer such as $1/3$ expressible as integers needs the scaling factor of 3 included. That points to the need to have some knowledge about the program answers *before* scaling the program, or to transform the whole program by using multipliers containing for example, all the denominators that can appear at run-time during the program execution.

Table 1b shows that, also for this benchmark, the additional precision increases the execution run-time of *clingo* by orders of magnitude.

5 Related work

Previous work translated *discrete* EC into ASP by reformulating the EC models as first-order stable models and translating the (almost universal) formulas of EC into a logic program that preserves stable models. Given a finite domain, EC2ASP (and its evolution, F2LP) compiles (discrete) EC formulas into ASP programs (Lee and Palla 2020; 2012). This translation scheme relies on two facts: the semantics of second-order circumscription and first-order stable models coincide on canonical formulas, and almost-universal formulas can be transformed into a logic program while preserving the stable models. As a result, computing models of EC descriptions can be done by computing the stable models of an appropriately generated program.

Clearly, approaches featuring discrete domains cannot faithfully handle continuous quantities such as time. In addition, because of their reliance on SAT solvers to find the stable models, they can only handle *safe* programs. In contrast, in the *s(CASP)* system, because of its direct support for predicates with arbitrary terms, constructive negation, and the novel *forall mechanism* (Marple *et al.* 2017; Arias *et al.* 2018), program safety is not a requirement. Thus, *s(CASP)* can model EC axioms much more directly, elegantly, and in continuous domains.

The approaches mentioned above assume discrete quantities and do not support reasoning about continuous time or change. As long as SAT-based ASP systems are used to model EC, continuous fluents cannot be straightforwardly expressed since they require unbound or dense domains for the variables. The work closest to incorporating continuous time makes use of SMT solvers. In this approach, constraints are incorporated into ASP and the grounded theory is executed using an SMT solver, as in Lee and Meng (2013). However, this approach has not been directly applied to modeling the EC. The closest tool chain is ASP Modulo Theory to SAT Modulo Theory (ASPMT2SMT) by Bartholomew and Lee (2014) that uses *gringo* to partially ground the ASPMT theories and generate constraints that are processed by *Z3* (de Moura and Bjørner 2008). However, regular, discrete ASP variables are at the heart of the model, and these are grounded and used to generate the constraints. Therefore, if these discrete variables approximate continuous variables in the model, the constraints generated will only approximate the conditions of the original problem and therefore their solutions will also be an approximation (or a subset) of the solutions for the real problem. In other words, the initial discretization done for the ASP variables will be propagated via the generated constraints to the final solutions that will, in the best case, be a discretized version of the actual solutions. As an example, if time is discretized, the solutions to the model will suffer from this discretization.

EC can be written as a (Horn-clause) logic program, but it cannot be executed directly by Prolog, as reported by Shanahan (2000), because it lacks some necessary features, such as constructive negation, deduction of negated atoms, and (to some extent) detection of infinite failure. The last one can to some degree be worked around by using variants of Prolog implementations that feature loop-breaking mechanisms in the presence of constraints, as done by Arias and Carro (2019). The other points may eventually need ad hoc coding for every example or crafting an interpreter able carry out deduction tailored to the task.

A common approach is to write a metainterpreter specific to the EC variant at hand. This can be as complex as writing a (specialized) theorem prover or, more often, a specialized interpreter whose correctness is difficult to ascertain (see the code by Chittaro and Montanari (1996)). Therefore, some Prolog implementations of EC do not completely formalize the calculus or implement a weaker version. In our case, we leverage the capabilities of *s(CASP)* to provide constructive, sound negation, negative rule heads, and loop detection.

6 Conclusions

We showed how EC can be modeled in *s(CASP)*, a goal-directed implementation of constraint ASP with predicates, with fewer limitations than other approaches. *s(CASP)*

can capture the notion of continuous time (and, in general, fluents) in EC thanks to its grounding-free top-down evaluation strategy. It can also represent complex models and answer queries in a flexible manner thanks to the use of constraints.

The main contribution of the paper is to show how EC can be directly modeled using s(CASP), an ASP system that seamlessly supports constraints. The modeling of EC using s(CASP) is more elegant and faithful to the original axioms compared to other approaches such as F2LP, where time has to be discretized. While other approaches such as ASPMT do support continuous domains, their reliance on SMT solvers makes constraints and dependencies among variables be lost during grounding. The use of s(CASP) brings other advantages: for example, in a query-driven system the trace of the proof / generation of the model is a justification for the answers to a query. Likewise, explanations for observations via abduction are also generated for free, thanks to the goal-directed, top-down execution of s(CASP).

To the best of the authors' knowledge, our approach is the only one that faithfully models continuous-time EC under the stable model semantics. All other approaches discretize time and thus do not model EC in a sound manner. Our approach supports both deduction and abduction with little or no additional effort.

The work reported in this paper can be seen as one of the first non-trivial applications of s(CASP). It illustrates the advantages that goal-directed ASP systems have over grounding and SAT solver-based ones for certain domains / classes of applications. The EC and its realization through s(CASP) is being used to model real-world avionics systems in the aerospace industry (Hall et al. 2021). Avionics systems are cyber-physical system consisting of sensors and actuators. Sensors are modeled as fluents while actuator actions correspond to events. The goal is to use the s(CASP) EC model to verify (timed) properties of these avionics systems as well as to identify gaps with respect to system requirements. Abductive capabilities of s(CASP) are being used to troubleshoot the system as well as to find gaps in system design. The overarching goal of the project is to build EC based tools for avionics software systems assurance.

Competing interests: The authors declare none.

Supplementary material

To view supplementary material for this article, please visit <http://dx.doi.org/10.1017/S1471068421000156>.

References

- ARIAS, J. AND CARRO, M. 2019. Description, implementation, and evaluation of a generic design for tabled CLP. *Theory and Practice of Logic Programming* 19, 3 (May), 412–448.
- ARIAS, J., CARRO, M., SALAZAR, E., MARPLE, K. AND GUPTA, G. 2018. Constraint Answer Set Programming without Grounding. *Theory and Practice of Logic Programming* 18, 3–4, 337–354.
- ARIAS, J., CHEN, Z., CARRO, M. AND GUPTA, G. 2019. Modeling and reasoning in event calculus using goal-directed constraint answer set programming. In *Pre-Proc. of the 29th Int'l. Symposium on Logic-based Program Synthesis and Transformation*.

- BARTHOLOMEW, M. AND LEE, J. 2014. System *aspmt2smt*: Computing ASPMT Theories by SMT Solvers. In *14th European Conference on Logics in Artificial Intelligence*. LNCS, vol. 8761. Springer, 529–542.
- CHITTARO, L. AND MONTANARI, A. 1996. Efficient Temporal Reasoning in the Cached Event Calculus. *Computational Intelligence* 12, 359–382.
- CLARK, K. L. 1978. Negation as Failure. In *Logic and Data Bases*, H. Gallaire and J. Minker, Eds. Springer, 293–322.
- DE MOURA, L. M. AND BJØRNER, N. 2008. Z3: An efficient SMT solver. In *Tools and Algorithms for the Construction and Analysis of Systems, 14th International Conference, TACAS 2008*, C. R. Ramakrishnan and J. Rehof, Eds. Lecture Notes in Computer Science, vol. 4963. Springer, 337–340.
- DOVIER, A., PONTELLI, E. AND ROSSI, G. 2000. A necessary condition for constructive negation in constraint logic programming. *Information Processing Letters* 74, 3–4, 147–156.
- ERDEM, E. AND LIFSCHITZ, V. 2003. Tight logic programs. *Theory and Practice of Logic Programming* 3, 4–5, 499–518.
- FERRARIS, P., LEE, J. AND LIFSCHITZ, V. 2011. Stable models and circumscription. *Artificial Intelligence* 175, 1, 236–263.
- GEBSER, M., KAMINSKI, R., KAUFMANN, B. AND SCHAUB, T. 2014. Clingo = ASP + Control: Preliminary Report. *arXiv 1405.3694*.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The stable model semantics for logic programming. In *5th International Conference on Logic Programming*, 1070–1080.
- GUPTA, G., BANSAL, A., MIN, R., SIMON, L. AND MALLYA, A. 2007. Coinductive Logic Programming and its Applications. In *23rd Int'l. Conference on Logic Programming*. Springer, 27–44.
- HALL, B., VARANASI, S. C., FIEDOR, J., ARIAS, J., BASU, K., LI, F., BHATT, D., DRISCOLL, K., SALAZAR, E. AND GUPTA, G. 2021. Knowledge-assisted reasoning of model-augmented system requirements with event calculus and goal-directed answer set programming. In *Proc. 8th Workshop on Horn Clause Verification and Synthesis*.
- HERMENEGILDO, M. V., BUENO, F., CARRO, M., LOPEZ-GARCIA, P., MERA, E., MORALES, J. AND PUEBLA, G. 2012. An overview of *ciao* and its design philosophy. *Theory and Practice of Logic Programming* 12, 1–2 (January), 219–252.
- HOLZBAUR, C. 1995. OFAI CLP(Q,R) Manual, Edition 1.3.3. Tech. Rep. TR-95-09, Austrian Research Institute for Artificial Intelligence, Vienna.
- KOWALSKI, R. AND SERGOT, M. 1989. A logic-based calculus of events. In *Foundations of Knowledge Base Management*. Springer, 23–55.
- LEE, J. AND MENG, Y. 2013. Answer set programming modulo theories and reasoning about continuous changes. In *23rd Int'l. Joint Conference on Artificial Intelligence*. 990–996.
- LEE, J. AND PALLA, R. 2012. Reformulating the situation calculus and the event calculus in the general theory of stable models and in answer set programming. *Journal of Artificial Intelligence Research* 43, 571–620.
- LEE, J. AND PALLA, R. 2020. F2LP: Computing answer sets of first order formulas. <https://github.com/azreasoners/F2LP>. Accessed on October, 2020.
- LIFSCHITZ, V. 1985. Computing circumscription. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence. Los Angeles, CA, USA, August 1985*. 121–127.
- LIFSCHITZ, V. 1999. *Action Languages, Answer Sets, and Planning*. Springer Berlin Heidelberg, 357–373.
- LIFSCHITZ, V., PEARCE, D. AND VALVERDE, A. 2001. Strongly Equivalent Logic Programs. *ACM Transactions on Computational Logic* 2, 4, 526–541.
- MAREK, V. W. AND TRUSZCZYNSKI, M. 1999. Stable Models and an Alternative Logic Programming Paradigm. In *The Logic Programming Paradigm: a 25-Year Perspective*. Springer-Verlag, 375–398.

- MARPLE, K., BANSAL, A., MIN, R. AND GUPTA, G. 2012. Goal-directed execution of answer set programs. In *Principles and Practice of Declarative Programming, PPDP'12, Leuven, Belgium - September 19 - 21, 2012*, D. D. Schreye, G. Janssens, and A. King, Eds. ACM, 35–44.
- MARPLE, K. AND GUPTA, G. 2014. Dynamic consistency checking in goal-directed answer set programming. *Theory and Practice of Logic Programming* 14, 4–5, 415–427.
- MARPLE, K., SALAZAR, E. AND GUPTA, G. 2017. Computing stable models of normal logic programs without grounding. *arXiv 1709.00501*.
- MARRIOTT, K. AND STUCKEY, P. J. 1998. *Programming with Constraints: An Introduction*. MIT Press.
- MARRIOTT, K., STUCKEY, P. J. AND WALLACE, M. 2006. Constraint logic programming. In *Foundations of Artificial Intelligence*. Vol. 2. Elsevier, 409–452.
- MCCARTHY, J. 1980. Circumscription - A form of non-monotonic reasoning. *Artificial Intelligence* 13, 1–2, 27–39.
- MELLARKOD, V. S., GELFOND, M. AND ZHANG, Y. 2008a. Integrating answer set programming and constraint logic programming. *Annals of Mathematics and Artificial Intelligence* 53, 1–4, 251–287.
- MELLARKOD, V. S., GELFOND, M. AND ZHANG, Y. 2008b. Integrating answer set programming and constraint logic programming. Tech. rep., Texas Tech University. October. This is a long version of (Mellarkod et al. 2008a), available at <https://www.depts.ttu.edu/cs/research/documents/46.pdf>.
- MUELLER, E. T. 2008a. Chapter 17: Event calculus. In *Handbook of Knowledge Representation*, F. van Harmelen, V. Lifschitz, and B. Porter, Eds. Foundations of AI, vol. 3. Elsevier, 671–708.
- MUELLER, E. T. 2008b. Discrete event calculus reasoner documentation. *Software documentation, IBM Thomas J. Watson Research Center, PO Box 704*. Available at: <http://decreasoner.sourceforge.net/>. Accessed on October, 2020.
- MUELLER, E. T. 2014. *Commonsense Reasoning: An Event Calculus Based Approach*. Morgan Kaufmann.
- SHANAHAN, M. 1999. The event calculus explained. In *Artificial Intelligence Today*. Springer, 409–430.
- SHANAHAN, M. 2000. An abductive event calculus planner. *The Journal of Logic Programming* 44, 1–3, 207–240.
- STUCKEY, P. 1991. Constructive negation for constraint logic programming. In *Proc. LICS'91*, 328–339.