

Relational Constraint Solving in SMT

Baoluo Meng¹, Andrew Reynolds¹, Cesare Tinelli¹, and Clark Barrett²

¹ Department of Computer Science, The University of Iowa

² Department of Computer Science, Stanford University

Abstract. Relational logic is useful for reasoning about computational problems with relational structures, including high-level system design, architectural configurations of network systems, ontologies, and verification of programs with linked data structures. We present a modular extension of an earlier calculus for the theory of finite sets to a theory of finite relations with such operations as transpose, product, join, and transitive closure. We implement this extension as a theory solver of the SMT solver CVC4. Combining this new solver with the finite model finding features of CVC4 enables several compelling use cases. For instance, native support for relations enables a natural mapping from Alloy, a declarative modeling language based on first-order relational logic, to SMT constraints. It also enables a natural encoding of several description logics with concrete domains, allowing the use of an SMT solver to analyze, for instance, Web Ontology Language (OWL) models. We provide an initial evaluation of our solver on a number of Alloy and OWL models which shows promising results.

Keywords: Relational Logic, SMT, Alloy, OWL

1 Introduction

Many computational problems require reasoning about relational structures. Examples include high-level system design, architectural configuration of network systems, reasoning about ontologies, and verification of programs with linked data structures. Relational logic is an appealing choice for reasoning about such problems. In this paper, we consider a many-sorted relational logic where relations of arity n are defined as sets of n -tuples with parametrized sorts for tuple elements. We define a version of this logic as a first-order theory of finite relations where relation terms are built from relation constants and variables, set operators, and relational operators such as join, transpose, product, and transitive closure.

In previous work [3], Bansal et al. developed a decision procedure for a theory of finite sets with cardinality constraints. The theory of finite relations presented here is an extension of that theory to relational constraints. We present a calculus for the satisfiability of quantifier-free formulas in this theory. Our calculus is in general refutation-sound and model-sound. It is also terminating and refutation-complete for a restricted class of quantifier-free formulas that has useful applications.

The calculus is explicitly designed to be implementable as a theory solver in SMT solvers based on the DPLL(T) architecture [14]. We have implemented a modular component for it in our SMT solver CVC4 [4], allowing CVC4 to solve constraints on relations over elements from any of the theories it supports. This relational extension of

CVC4’s native input language enables natural mappings to SMT formulas from several modeling languages based on relations. This includes Alloy, a formal language based on first-order relational logic, as well as ontology languages such as OWL. A significant potential advantage of these mappings is that they bring to these languages the power of SMT solvers to reason natively about a variety of interpreted types, something that is challenging for existing reasoners for these languages.

1.1 Related Work

Alloy is a well-known declarative modeling language based on a relational logic with transitive closure, set cardinality, and integer arithmetic operators [13]. Alloy specifications, or *models*, can be analyzed for consistency or other entailment properties with the Alloy Analyzer, a static analyzer based on encoding models as propositional logic formulas, and passing them to an off-the-shelf propositional satisfiability (SAT) solver. This approach limits the analysis to models with explicit and concrete cardinality bounds on the relations involved; hence it is appropriate only for proving the consistency of a model or for disproving that a given property, encoded as a formula, holds for a model. Despite these limitations, Alloy and its analyzer have been quite useful for lightweight modeling and analysis of software systems. An earlier attempt to solve Alloy constraints without artificial cardinality bounds on relations was made by Ghazi and Taghdiri [8] using SMT solvers. They developed a translation from a subset of Alloy’s language to the SMT-LIB language [5] and used the SMT solver Yices [7] to solve the resulting constraints. That approach can prove some properties of certain Alloy models, but it still requires explicitly *finitizing* relations when dealing with transitive closure, limiting the kind of properties that can be proven in models that contain applications of transitive closure. Later, the same authors introduced more general methods [9, 10], implemented in the AlloyPE tool, to axiomatize relational operators as SMT formulas without finitization while covering the entire core language of Alloy. However, since quantified relational logic is in general undecidable and quantifiers are heavily used in the translation, the resulting SMT formulas translated from Alloy are often difficult or even impossible for SMT solvers to solve, especially when transitive closure is involved.

Description Logics (DLs) [1, 2] are decidable fragments of relational logic explicitly developed for efficient knowledge representation and reasoning. They consider on purpose only unary and binary relations. The main building blocks of DLs are individuals, concepts, roles as well as operations over these, where concepts represent sets of individuals, roles represent binary relations between individuals, and operations include membership, subset, relational composition (join) and equality. Restricted use of quantifiers is allowed in DLs to encode more expressive constraints on roles and concepts. OWL [20], a standardized semantic web ontology language, represents an important application of DLs to ontological modeling. It consists of entities similar to those in DLs except for superficial differences in concrete syntax and for the inclusion of additional features that make reasoning about OWL models undecidable in general. Many efficient reasoners have been built for reasoning about OWL ontologies written in restricted fragments of OWL. These include Konclude [16], FaCT++ [18], and Chainsaw [19].

1.2 Formal Preliminaries

We define our theory of relations and our calculus in the context of many-sorted first-order logic with equality. We assume the reader is familiar with the following notions from that logic: signature, term, literal, formula, free variable, interpretation, and satisfiability of a formula in an interpretation (see, e.g., [6] for more details). Let Σ be a many-sorted signature. We will use \approx as the (infix) logical symbol for equality—which has type $\sigma \times \sigma$ for all sorts σ in Σ and is always interpreted as the identity relation. We assume all signatures Σ contain the Boolean sort `Bool`, always interpreted as the binary set $\{\text{true}, \text{false}\}$, and a Boolean constant symbol `true` for `true`. Without loss of generality, we assume \approx is the only predicate symbol in Σ , as all other predicates may be modeled as functions with return sort `Bool`. We will commonly write, e.g. $P(x)$, as shorthand for $P(x) \approx \text{true}$ where $P(x)$ has sort `Bool`. We write $s \not\approx t$ as an abbreviation of $\neg s \approx t$. If e is a term or a formula, we denote by $\text{Vars}(e)$ the set of e 's free variables, extending the notation to tuples and sets of terms or formulas as expected. We write $\varphi[x]$ to indicate that all the free variables of a formula φ are from tuple x .

If φ is a Σ -formula and \mathcal{I} a Σ -interpretation, we write $\mathcal{I} \models \varphi$ if \mathcal{I} satisfies φ . If t is a term, we denote by $t^{\mathcal{I}}$ the value of t in \mathcal{I} . A *theory* is a pair $T = (\Sigma, \mathbf{I})$, where Σ is a signature and \mathbf{I} is a class of Σ -interpretations that is closed under variable reassignment (i.e., every Σ -interpretation that differs from one in \mathbf{I} only in how it interprets the variables is also in \mathbf{I}). \mathbf{I} is also referred to as the *models* of T . A Σ -formula φ is *satisfiable* (resp., *unsatisfiable*) in T if it is satisfied by some (resp., no) interpretation in \mathbf{I} . A set Γ of Σ -formulas *entails in T* a Σ -formula φ , written $\Gamma \models_T \varphi$, if every interpretation in \mathbf{I} that satisfies all formulas in Γ satisfies φ as well. We write $\models_T \varphi$ as an abbreviation for $\emptyset \models_T \varphi$. We write $\Gamma \models \varphi$ to denote that Γ entails φ in the class of all Σ -interpretations. Two Σ -formulas are *equisatisfiable in T* if for every model \mathcal{A} of T that satisfies one, there is a model of T that satisfies the other and differs from \mathcal{A} at most over the free variables not shared by the two formulas. When convenient, we will treat a finite set of formulas as the conjunction of its elements and vice versa.

2 A Relational Extension to a Theory of Finite Sets

A many-sorted theory of finite sets with cardinality T_S is described in detail in our previous work [3]. The theory T_S includes a set sort constructor $\text{Set}(\alpha)$ parametrized by the sort of the set elements. The theory T_S can be combined with any other theory T Nelson-Oppen-style, by instantiating the parameter α with any sort in T . The signature of theory T_S includes function and predicate symbols for set union (\sqcup), intersection (\sqcap), difference (\setminus), the empty set ($\{\}$), singleton set construction ($\{_ \}$), set inclusion (\sqsubseteq), and membership (\in), all interpreted as expected. A sound, complete and terminating tableaux-style calculus for the theory T_S is implemented in the CVC4 SMT solver [4].

In this section, we describe an extension T_R of this theory which, however, does not include a set cardinality operator or cardinality constraints.³ The new theory T_R extends T_S with a parametric tuple sort $\text{Tuple}_n(\alpha_1, \dots, \alpha_n)$ for every $n > 0$ and various relational operators defined over sets of tuples, that is, over values whose sort is an instance

³ A further extension of the theory to cardinality constraints is planned for future work.

Set symbols:

$$\begin{aligned} [] : \text{Set}(\alpha) & \quad \sqcup, \sqcap, \setminus : \text{Set}(\alpha) \times \text{Set}(\alpha) \rightarrow \text{Set}(\alpha) & \in : \alpha \times \text{Set}(\alpha) \rightarrow \text{Bool} \\ [-] : \alpha \rightarrow \text{Set}(\alpha) & \quad \sqsubseteq : \text{Set}(\alpha) \times \text{Set}(\alpha) \rightarrow \text{Bool} \end{aligned}$$

Relation symbols:

$$\begin{aligned} \langle -, \dots, - \rangle : \alpha_1 \times \dots \times \alpha_n & \rightarrow \text{Tup}_n(\alpha_1, \dots, \alpha_n) \\ * : \text{Rel}_m(\alpha) \times \text{Rel}_n(\beta) & \rightarrow \text{Rel}_{m+n}(\alpha, \beta) \\ \bowtie : \text{Rel}_{p+1}(\alpha, \gamma) \times \text{Rel}_{q+1}(\gamma, \beta) & \rightarrow \text{Rel}_{p+q}(\alpha, \beta) \text{ with } p+q > 0 \\ -^{-1} : \text{Rel}_m(\alpha_1, \dots, \alpha_m) & \rightarrow \text{Rel}_m(\alpha_m, \dots, \alpha_1) & -^+ : \text{Rel}_2(\alpha, \alpha) \rightarrow \text{Rel}_2(\alpha, \alpha) \end{aligned}$$

where $m, n > 0$, $\alpha = (\alpha_1, \dots, \alpha_m)$, $\beta = (\beta_1, \dots, \beta_n)$, and $\text{Rel}_n(\gamma) = \text{Set}(\text{Tup}_n(\gamma))$.

Fig. 1: Signature Σ_R of our relational theory T_R .

of $\text{Set}(\text{Tup}_n(\alpha_1, \dots, \alpha_n))$. We call any sort σ of the form $\text{Set}(\text{Tup}_n(\sigma_1, \dots, \sigma_n))$ a *relational sort (of arity n)* and abbreviate it as $\text{Rel}_n(\sigma_1, \dots, \sigma_n)$.

The full signature Σ_R of T_R is defined in Figure 1. Note that the function symbols $*$, \bowtie , and $-^{-1}$ are not only parametric but also overloaded, as they apply to relational sorts $\text{Rel}_k(\sigma)$ of different arities k . The models of T_R are the expansions of the models of T_S that interpret $\langle -, \dots, - \rangle$ as the n -tuple constructor, $*$ as relational product, \bowtie as relational join, $-^{-1}$ as the transpose operator, and $-^+$ as the transitive closure operator. A *relation term* is a Σ_R -term of some relational sort. A *tuple term* is a Σ_R -term of some tuple sort. A T_R -*constraint* is a (dis)equality of the form $(\neg)s \approx t$, where s and t are Σ_R -terms. We write $s \not\approx t$ and $[t_1, \dots, t_n]$ with $n > 1$ as an abbreviation of $\neg(s \in t \approx \text{true})$ and $[t_1] \sqcup \dots \sqcup [t_n]$, respectively.

3 A Calculus for the Relational Extension

In this section, we describe a tableaux-style calculus for determining the satisfiability of finite sets of T_R -constraints. The calculus consists of a set of derivation rules similar to those in the calculus from [3] that deal with set constraints as well as new rules to process T_R -constraints. For simplicity, we will implicitly assume that the sort of any set or relation term is flat (i.e., set or relation elements are not themselves sets or relations) and allow only variables as terms of element sorts. Nested sets and relations and more complex element terms can be processed in a standard way by using a Nelson-Oppen-style approach which we will not discuss here.

The derivation rules modify a *state* data structure, where a state is either the distinguished state *unsat* or a set \mathcal{S} of T_R -constraints. The rules are provided in Figure 2 and 3 in *guarded assignment form*. In such form, the premises of a rule refer to the current state \mathcal{S} and the conclusion describes how \mathcal{S} is changed by the rule's application. Rules with two or more conclusions, separated by the symbol \parallel , are non-deterministic branching rules. In the rules, we write \mathcal{S}, c as an abbreviation of $\mathcal{S} \cup \{c\}$, and denote by $\mathcal{T}(\mathcal{S})$ the set of all terms and subterms occurring in \mathcal{S} . We define the following closure

$$\begin{array}{c}
\text{INTER UP} \\
\frac{x \in s \in \mathcal{S}^* \quad x \in t \in \mathcal{S}^* \quad s \sqcap t \in \mathcal{T}(\mathcal{S})}{\mathcal{S} := \mathcal{S}, x \in s \sqcap t} \\
\\
\text{INTER DOWN} \\
\frac{x \in s \sqcap t \in \mathcal{S}^*}{\mathcal{S} := \mathcal{S}, x \in s, x \in t} \\
\\
\text{UNION UP} \\
\frac{x \in u \in \mathcal{S}^* \quad u \in \{s, t\} \quad s \sqcup t \in \mathcal{T}(\mathcal{S})}{\mathcal{S} := \mathcal{S}, x \in s \sqcup t} \\
\\
\text{UNION DOWN} \\
\frac{x \in s \sqcup t \in \mathcal{S}^*}{\mathcal{S} := \mathcal{S}, x \in s \parallel \mathcal{S} := \mathcal{S}, x \in t} \\
\\
\text{DIFF UP} \\
\frac{x \in s \in \mathcal{S}^* \quad s \setminus t \in \mathcal{T}(\mathcal{S})}{\mathcal{S} := \mathcal{S}, x \in t \parallel \mathcal{S} := \mathcal{S}, x \in s \setminus t} \\
\\
\text{DIFF DOWN} \\
\frac{x \in s \setminus t \in \mathcal{S}^*}{\mathcal{S} := \mathcal{S}, x \in s, x \notin t} \\
\\
\text{SINGLE UP} \quad \text{SINGLE DOWN} \quad \text{EMPTY UNSAT} \\
\frac{[x] \in \mathcal{T}(\mathcal{S})}{\mathcal{S} := \mathcal{S}, x \in [x]} \quad \frac{x \in [y] \in \mathcal{S}^*}{\mathcal{S} := \mathcal{S}, x \approx y} \quad \frac{x \in [] \in \mathcal{S}^*}{\text{unsat}} \\
\\
\text{SET DISEQ} \quad \text{EQ UNSAT} \\
\frac{s \not\approx t \in \mathcal{S}^*}{\mathcal{S} := \mathcal{S}, z \in s, z \notin t \parallel \mathcal{S} := \mathcal{S}, z \notin s, z \in t} \quad \frac{(t \not\approx t) \in \mathcal{S}^*}{\text{unsat}}
\end{array}$$

Fig. 2: Basic rules for set intersection, union, difference, singleton, disequality and contradiction. In SET DISEQ, z is a fresh variable.

operator for \mathcal{S} where \models_{tup} denotes entailment in the Σ_R -theory of tuples:⁴

$$\begin{aligned}
\mathcal{S}^* = & \{s \approx t \mid s, t \in \mathcal{T}(\mathcal{S}), \mathcal{S} \models_{\text{tup}} s \approx t\} \cup \\
& \{s \not\approx t \mid s, t \in \mathcal{T}(\mathcal{S}), \mathcal{S} \models_{\text{tup}} s \approx s' \wedge t \approx t' \text{ for some } s' \not\approx t' \in \mathcal{S}\} \cup \\
& \{s \in t \mid s, t \in \mathcal{T}(\mathcal{S}), \mathcal{S} \models_{\text{tup}} s \approx s' \wedge t \approx t' \text{ for some } s' \in t' \in \mathcal{S}\}
\end{aligned}$$

The set \mathcal{S}^* is computable by extending congruence closure procedures with a rule for deducing consequences of tuple equalities of the form $\langle s_1, \dots, s_n \rangle \approx \langle t_1, \dots, t_n \rangle$.

A derivation rule *applies* to a state \mathcal{S} if all the conditions in the rule's premises hold for \mathcal{S} *and* the rule application is not redundant. An application of a rule to a state \mathcal{S} with a conclusion $\mathcal{S} \cup \{\varphi_1[\mathbf{x}_1, \mathbf{z}], \dots, \varphi_n[\mathbf{x}_n, \mathbf{z}]\}$, where \mathbf{z} are the fresh variables introduced by the rule's application (if any), is *redundant* if \mathcal{S} already contains $\varphi_1[\mathbf{x}_1, \mathbf{t}], \dots, \varphi_n[\mathbf{x}_n, \mathbf{t}]$ for some terms \mathbf{t} .

For simplicity and without loss of generality, we consider only initial states \mathcal{S}_0 that contain no variables of tuple sorts $\text{Tuple}_n(\sigma_1, \dots, \sigma_n)$, since such variables can be replaced by a tuple $\langle x_1, \dots, x_n \rangle$ where each x_i is a variable of sort σ_i . We also assume that \mathcal{S}_0 contains no atoms of the form $s \sqsubseteq t$, since they can be replaced by $s \approx s \sqcap t$, or disequalities $\langle s_1, \dots, s_n \rangle \not\approx \langle t_1, \dots, t_n \rangle$ between tuple terms, since those can be treated by guessing a disequality $s_i \not\approx t_i$ between two of their respective components. All derivation rules preserve these restrictions.

Figure 2 presents the basic rules for the core set constraints in our theory. For each set operator, Figure 2 contains a *downward* rule and an *upward* rule. Given a mem-

⁴ Note that this theory has all the function symbols of Σ_R , not just the tuple constructors $\langle -, \dots, - \rangle$. The extra symbols are treated as *uninterpreted*.

$$\begin{array}{c}
\text{TRANSP UP} \\
\frac{\langle x_1, \dots, x_n \rangle \in R \in \mathcal{S}^* \quad R^{-1} \in \mathcal{T}(\mathcal{S})}{\mathcal{S} := \mathcal{S}, \langle x_n, \dots, x_1 \rangle \in R^{-1}} \\
\\
\text{TRANSP DOWN} \\
\frac{\langle x_1, \dots, x_n \rangle \in R^{-1} \in \mathcal{S}^*}{\mathcal{S} := \mathcal{S}, \langle x_n, \dots, x_1 \rangle \in R} \\
\\
\text{PROD UP} \\
\frac{\langle x_1, \dots, x_m \rangle \in R_1 \in \mathcal{S}^* \quad \langle y_1, \dots, y_n \rangle \in R_2 \in \mathcal{S}^* \quad R_1 * R_2 \in \mathcal{T}(\mathcal{S})}{\mathcal{S} := \mathcal{S}, \langle x_1, \dots, x_m, y_1, \dots, y_n \rangle \in R_1 * R_2} \\
\\
\text{PROD DOWN} \\
\frac{\langle x_1, \dots, x_m, y_1, \dots, y_n \rangle \in R_1 * R_2 \in \mathcal{S}^* \quad \text{ar}(R_1) = m}{\mathcal{S} := \mathcal{S}, \langle x_1, \dots, x_m \rangle \in R_1, \langle y_1, \dots, y_n \rangle \in R_2} \\
\\
\text{JOIN UP} \\
\frac{\langle x_1, \dots, x_m, z \rangle \in R_1, \langle z, y_1, \dots, y_n \rangle \in R_2 \in \mathcal{S}^* \quad m + n > 0 \quad R_1 \bowtie R_2 \in \mathcal{T}(\mathcal{S})}{\mathcal{S} := \mathcal{S}, \langle x_1, \dots, x_m, y_1, \dots, y_n \rangle \in R_1 \bowtie R_2} \\
\\
\text{JOIN DOWN} \\
\frac{\langle x_1, \dots, x_m, y_1, \dots, y_n \rangle \in R_1 \bowtie R_2 \in \mathcal{S}^* \quad \text{ar}(R_1) = m + 1}{\mathcal{S} := \mathcal{S}, \langle x_1, \dots, x_m, z \rangle \in R_1, \langle z, y_1, \dots, y_n \rangle \in R_2} \\
\\
\text{TCLOS UP I} \\
\frac{\langle x_1, x_2 \rangle \in R \in \mathcal{S}^* \quad R^+ \in \mathcal{T}(\mathcal{S})}{\mathcal{S} := \mathcal{S}, \langle x_1, x_2 \rangle \in R^+} \\
\text{TCLOS UP II} \\
\frac{\langle x_1, x_2 \rangle \in R^+, \langle x_2, x_3 \rangle \in R^+ \in \mathcal{S}^*}{\mathcal{S} := \mathcal{S}, \langle x_1, x_3 \rangle \in R^+} \\
\\
\text{TCLOS DOWN} \\
\frac{\langle x_1, x_2 \rangle \in R^+ \in \mathcal{S}^*}{\mathcal{S} := \mathcal{S}, \langle x_1, x_2 \rangle \in R \quad \parallel \quad \mathcal{S} := \mathcal{S}, \langle x_1, z \rangle \in R, \langle z, x_2 \rangle \in R \\
\parallel \quad \mathcal{S} := \mathcal{S}, \langle x_1, z_1 \rangle \in R, \langle z_1, z_2 \rangle \in R^+, \langle z_2, x_2 \rangle \in R, z_1 \not\approx z_2}
\end{array}$$

Fig. 3: Basic relational derivation rules. Letters z, z_1, z_2 denote fresh variables.

bership constraint $x \in s$, the downward rules infer either additional membership constraints over the immediate subterms of s , or an equality in the case where s is $\{y\}$. For example, rule INTER DOWN, infers the constraints $x \in s$ and $x \in t$ if \mathcal{S}^* contains the constraint $x \in s \sqcap t$. The upward rules handle the case where some set s occurs in \mathcal{S} , and infer membership constraints of the form $x \in s$ based on other constraints from \mathcal{S} . Rule SET DISEQ introduces a witness for a disequality between two sets s and t by using a fresh variable z to assert that there is an element that is in s but not t , or in t but not in s . There are two rules for deriving unsat from trivially unsatisfiable constraints in \mathcal{S} : membership constraints of the form $x \in \emptyset$ (EMPTY UNSAT) and disequalities of the form $t \not\approx t$ (EQ UNSAT).

We supplement the set-specific rules with an additional set of rules for T_R -constraints, given in Figure 3. From the membership of a tuple in the transpose of a relation R , rule TRANSP DOWN concludes that the reverse of the tuple is in R . Conversely, rule TRANSP UP ensures that the reverse of a tuple is in the transpose of a relation R if the tuple is in R and R^{-1} occurs in \mathcal{S} . From the constraint that a tuple t belongs to the join of two relations R_1 and R_2 with arities m and n respectively, rule JOIN DOWN infers that R_1 contains a tuple t_1 whose last element (named using a fresh Skolem

variable z) is the first element of a tuple t_2 in R_2 , where t is the join of t_1 and t_2 . The JOIN UP rule computes the join of pairs of tuples explicitly asserted to belong to a relation R_1 and a relation R_2 , respectively, provided that $R_1 \bowtie R_2$ is a term in \mathcal{S} . The PROD DOWN and PROD UP rules are defined similarly for the product of relations. The rules TCLOS UP I and TCLOS UP II compute members of the transitive closure of R based on the (currently asserted) members of R . When it can be inferred that a tuple $\langle x_1, x_2 \rangle$ belongs to the transitive closure of a binary relation R , TCLOS DOWN can produce three alternative conclusions. In reachability terms, the first conclusion considers the case that x_2 is directly reachable from x_1 in the graph induced by R , the second that x_2 is reachable from x_1 in two steps, and the third that it is reachable in more steps. Note that the third case may lead to additional applications of TCLOS DOWN, possibly indefinitely, if the other constraints in \mathcal{S} (implicitly) entail that $\langle x_1, x_2 \rangle$ does not in fact belong to R .

Example 1. Let $\mathcal{S} = \{\langle x, y \rangle \in R^{-1}, R \approx S, \langle y, x \rangle \notin S\}$. By rule TRANSP DOWN, we can derive a constraint $\langle y, x \rangle \in R$, leading to a new \mathcal{S} : $\{\langle x, y \rangle \in R^{-1}, R \approx S, \langle y, x \rangle \notin S, \langle y, x \rangle \in R\}$. Then, $\langle y, x \rangle \in R$ is both equal and disequal to true in \mathcal{S}^* . Thus, we can derive unsat by EQ UNSAT, and conclude that \mathcal{S} is T_R -unsatisfiable. \square

Example 2. Let \mathcal{S} be $\{\langle x \rangle \in R, \langle y \rangle \in R, R * R \approx S \sqcap T, \langle y, x \rangle \notin T\}$. By rule PROD UP, we derive constraints $\langle x, y \rangle \in R * R, \langle y, x \rangle \in R * R, \langle x, x \rangle \in R * R$ and $\langle y, y \rangle \in R * R$. By set reasoning rule INTER DOWN, we derive another four constraints $\langle x, y \rangle \in S, \langle y, x \rangle \in S, \langle x, y \rangle \in T$, and $\langle y, x \rangle \in T$, leading to a contradiction with $\langle y, x \rangle \notin T$. Thus, we can derive unsat by rule EQ UNSAT. \square

Example 3. Let \mathcal{S} be $\{\langle x, y \rangle \in R, \langle y, z \rangle \in R, \langle x, z \rangle \notin R^+\}$. By rule TCLOS UP I, we derive two new constraints $\langle x, y \rangle \in R^+$ and $\langle y, z \rangle \in R^+$. Then, we can derive another constraint $\langle x, z \rangle \in R^+$, by rule TCLOS UP II, in contradiction with $\langle x, z \rangle \notin R^+$. Thus, we can derive unsat by rule EQ UNSAT. \square

Example 4. Let \mathcal{S} be $\{\langle x, y \rangle \in R^+, \langle x, y \rangle \notin R\}$. By rule TCLOS DOWN, we construct a derivation tree with three child branches, which add to \mathcal{S} the sets $\{\langle x, y \rangle \in R\}$, $\{\langle x, z \rangle \in R, \langle z, y \rangle \in R\}$, and $\{\langle x, z_1 \rangle \in R, \langle z_1, z_2 \rangle \in R^+, \langle z_2, y \rangle \in R, z_1 \not\approx z_2\}$ respectively, where z_1 and z_2 are fresh variables. By rule EQ UNSAT, we can derive unsat in the first branch. Since no rules apply to the second branch, we can conclude, as we will see, that \mathcal{S} is T_R -satisfiable. \square

4 Calculus Correctness

In this section, we formalize the correctness properties satisfied by our calculus. These include refutation and model soundness in general and termination over a fragment of our language of constraints. The rules of the calculus define a notion of *derivation trees*. These are possibly infinite trees whose nodes are states where the children of each non-leaf node are the result of applying one of the derivation rules of the calculus to that node. A finite branch of a derivation tree is *closed* if it ends with unsat; it is *saturated* if no rules apply to its leaf. A derivation tree is *closed* if all of its branches are closed.

Lemma 1. *For all applications of a rule from the calculus, the parent state \mathcal{S} is T_R -satisfiable only if there exists a child \mathcal{S}' that is T_R -satisfiable.*

Proof. We show this holds for each rule of the calculus. First, we remark that by construction, \mathcal{S}^* is entailed by \mathcal{S} in T_R . For the rules EMPTY UNSAT and EQ UNSAT with conclusion *unsat*, we have that \mathcal{S} is clearly T_R -unsatisfiable, since \mathcal{S}^* contains a constraint that is T_R -unsatisfiable.

For the rules UNION DOWN and DIFF UP, assuming \mathcal{S} is satisfied by some interpretation \mathcal{I} , that interpretation also satisfies one of the child branches. For rule SET DISEQ, if \mathcal{S} is T_R -satisfiable and $s \not\approx t \in \mathcal{S}^*$, then all models of \mathcal{S} are such that $\exists x.x \notin u \wedge x \in v$ for $\{u, v\} = \{s, t\}$, and thus one of the child branches is T_R -satisfiable. Similarly, for rule TCLOS DOWN, if $\langle x_1, x_2 \rangle \in R^+ \in \mathcal{S}^*$, then if \mathcal{S} is satisfied by some interpretation \mathcal{I} , that interpretation must satisfy either $\langle x_1, x_2 \rangle \in R$ or $\langle x_1, y_1 \rangle, \langle y_1, y_2 \rangle, \dots, \langle y_{n-1}, y_n \rangle, \langle y_n, x_2 \rangle \in R$ for some y_1, \dots, y_n where $n \geq 1$ and $y_i^{\mathcal{I}} \neq y_j^{\mathcal{I}}$ for $i \neq j$. We have that \mathcal{I} satisfies the first conclusion of TCLOS DOWN in the former case, the second conclusion of TCLOS DOWN in the latter case where $n = 1$ and $z^{\mathcal{I}} = y_1^{\mathcal{I}}$, and the third conclusion of TCLOS DOWN in the latter case where $n > 1$, $z_1^{\mathcal{I}} = y_1^{\mathcal{I}}$ and $z_2^{\mathcal{I}} = y_n^{\mathcal{I}}$. Thus, at least one child branch is T_R -satisfiable.

All other rules of the calculus have a single conclusion of the form $\mathcal{S}' = \mathcal{S} \cup \{\varphi_1, \dots, \varphi_n\}$. For each of these rules, it can be shown that $\mathcal{S} \models_{T_R} \exists z \varphi_1 \wedge \dots \wedge \varphi_n$, where z are the fresh variables introduced by the rule application. For example, for INTER DOWN, we have that $x \in s \sqcap t \in \mathcal{S}^*$. We have that $\mathcal{S}' = \mathcal{S} \cup \{x \in s, x \in t\}$. Since $x \in s \sqcap t \in \mathcal{S}^*$, we have that $\mathcal{S} \models_{T_R} x \in s \sqcap t$, and hence $\mathcal{S} \models_{T_R} x \in s, x \in t$ as well. Thus \mathcal{S}' and \mathcal{S} are equivalent for all rule applications with a single conclusion, and the statement holds for these rules as well. \square

Proposition 1 (Refutation Soundness). *If there is a closed derivation tree with root node \mathcal{S} , then \mathcal{S} is T_R -unsatisfiable.*

Proof. Assume there exists a closed derivation tree with root node \mathcal{S} . Since all leaves of this tree are *unsat*, then for each node \mathcal{S}' of this derivation tree, it follows from Lemma 1 by induction on the structure of the derivation tree that \mathcal{S}' is T_R -unsatisfiable. Thus, \mathcal{S} is T_R -unsatisfiable and the proposition holds. \square

Proposition 2 (Model Soundness). *Let \mathcal{S} be the leaf of a saturated branch in a derivation tree with root node \mathcal{S}_0 . There is a model \mathcal{I} of T_R that satisfies \mathcal{S}_0 where:*

1. for all $S \in \text{Vars}(\mathcal{S})$ of set sort, $S^{\mathcal{I}} = \{x^{\mathcal{I}} \mid x \in S \in \mathcal{S}^*\}$, and
2. for all other $x, y \in \text{Vars}(\mathcal{S})$, $x^{\mathcal{I}} = y^{\mathcal{I}}$ if and only if $x \approx y \in \mathcal{S}^*$.

Proof. Let \mathcal{S} be the leaf of a saturated branch in a derivation tree with root node \mathcal{S}_0 . Let \mathcal{I} be a model of T_R with the above properties.

We argue that \mathcal{I} is a model of \mathcal{S} . We first show an intermediate fact involving the interpretation of set terms in \mathcal{I} . For any set term s , we define:

$$\text{Elements}(s) = \{x^{\mathcal{I}} \mid x \in s \in \mathcal{S}^*\}$$

We show by structural induction that for all sets $s \in \mathcal{T}(\mathcal{S})$:

$$\text{Elements}(s) = s^{\mathcal{I}} \tag{1}$$

We show this for each case of s .

Case 1 (s is a variable). In this case, $\text{Elements}(s) = s^{\mathcal{I}}$ by the first part of our construction of \mathcal{I} .

Case 2 (s is $[\]$). Since **EMPTY UNSAT** does not apply to \mathcal{S} , there are no constraints of the form $x \in s$ in \mathcal{S}^* , thus $\text{Elements}(s) = \emptyset = s^{\mathcal{I}}$.

Case 3 (s is $[x]$). Since **SINGLE UP** does not apply to \mathcal{S} , there must be a constraint of the form $x \in s$ in \mathcal{S}^* , thus $s^{\mathcal{I}} \subseteq \text{Elements}(s)$. For all constraints of the form $y \in s \in \mathcal{S}^*$, since **SINGLE DOWN** does not apply to \mathcal{S} , we have that $y \approx x \in \mathcal{S}$, and thus $y^{\mathcal{I}} = x^{\mathcal{I}}$. Thus, $\text{Elements}(s) \subseteq s^{\mathcal{I}}$, and thus $s^{\mathcal{I}} = \text{Elements}(s)$.

Case 4 (s is s_1^{-1}). By the induction hypothesis we have that $\text{Elements}(s_1) = s_1^{\mathcal{I}}$, and $s^{\mathcal{I}} = s_1^{-1\mathcal{I}} = \text{Elements}(s_1)^{-1}$. Since **TRANSP UP** does not apply to \mathcal{S} , for each $\langle x_1, \dots, x_n \rangle \in s_1 \in \mathcal{S}^*$, we have that $\langle x_n, \dots, x_1 \rangle \in s_1^{-1} \in \mathcal{S}$. Thus, for each $\langle x_1, \dots, x_n \rangle^{\mathcal{I}} \in \text{Elements}(s_1)$, we have that $(\langle x_1, \dots, x_n \rangle^{\mathcal{I}})^{-1} \in \text{Elements}(s_1^{-1})$, and thus we have that $\text{Elements}(s_1)^{-1} \subseteq \text{Elements}(s)$. Since **TRANSP DOWN** does not apply, for each $\langle x_1, \dots, x_n \rangle \in s_1^{-1} \in \mathcal{S}^*$, we have that $\langle x_n, \dots, x_1 \rangle \in s_1 \in \mathcal{S}$. Thus, we have that $\text{Elements}(s) \subseteq \text{Elements}(s_1)^{-1}$, thus $s^{\mathcal{I}} = \text{Elements}(s_1)^{-1} = \text{Elements}(s)$.

Case 5 (s is s_1^+). By the induction hypothesis we have that $\text{Elements}(s_1) = s_1^{\mathcal{I}}$, and $s^{\mathcal{I}} = s_1^{+\mathcal{I}} = \text{Elements}(s_1)^+$. Since **TCLOS UP I** does not apply to \mathcal{S} , we have that for each $\langle x_1, x_2 \rangle \in s_1 \in \mathcal{S}^*$, we have that $\langle x_1, x_2 \rangle \in s_1^+ \in \mathcal{S}$. Thus, we have that $\text{Elements}(s_1) \subseteq \text{Elements}(s)$. Since **TCLOS UP II** does not apply to \mathcal{S} , we have that $\text{Elements}(s)^+ = \text{Elements}(s)$, and hence by the previous subset relation we have that $\text{Elements}(s_1)^+ \subseteq \text{Elements}(s)$.

To show $\text{Elements}(s) \subseteq \text{Elements}(s_1)^+$, we assign a measure $m(\langle x_1, x_2 \rangle)$ to each tuple $\langle x_1, x_2 \rangle$ for which $\langle x_1, x_2 \rangle \in s_1^+ \in \mathcal{S}^*$. For each such tuple, since **TCLOS DOWN** does not apply to \mathcal{S} , we are in one of three cases. If \mathcal{S} is in a branch beneath either the first or second conclusion of **TCLOS DOWN** applied to $\langle x_1, x_2 \rangle \in s_1^+ \in \mathcal{S}^*$, we say $m(\langle x_1, x_2 \rangle) = 0$. Otherwise, if \mathcal{S} is in a branch beneath the third conclusion, we have that $\langle z_1, z_2 \rangle \in s_1^+ \in \mathcal{S}$ for some variables z_1, z_2 , and we say $m(\langle x_1, x_2 \rangle) = m(\langle z_1, z_2 \rangle) + 1$. Notice that $m(\langle x_1, x_2 \rangle)$ is finite for all $\langle x_1, x_2 \rangle$ since \mathcal{S} is a finite set. By induction on $m(\langle x_1, x_2 \rangle)$, we show that $\langle x_1, x_2 \rangle^{\mathcal{I}} \in \text{Elements}(s_1)^+$. If $m(\langle x_1, x_2 \rangle) = 0$, then either $\langle x_1, x_2 \rangle \in s_1$ or for some z both $\langle x_1, z \rangle \in s_1$ and $\langle z, x_2 \rangle \in s_1$ are in \mathcal{S} . In both cases, we have that $\langle x_1, x_2 \rangle^{\mathcal{I}} \in \text{Elements}(s_1)^+$. If $m(\langle x_1, x_2 \rangle) > 0$, then it must be the case that $\langle x_1, z_1 \rangle \in s_1$, $\langle z_1, z_2 \rangle \in s_1^+$, and $\langle z_2, x_2 \rangle \in s_1$ are in \mathcal{S} for some z_1, z_2 . Since $m(\langle z_1, z_2 \rangle) = m(\langle x_1, x_2 \rangle) - 1$, by the induction hypothesis, we have that $\langle z_1, z_2 \rangle^{\mathcal{I}} \in \text{Elements}(s_1)^+$. The other two constraints in \mathcal{S} imply that $\langle x_1, z_1 \rangle^{\mathcal{I}}$ and $\langle z_2, x_2 \rangle^{\mathcal{I}}$ are in $\text{Elements}(s_1)$ and hence $\langle x_1, x_2 \rangle^{\mathcal{I}} \in \text{Elements}(s_1)^+$. Thus, we have that $\text{Elements}(s) \subseteq \text{Elements}(s_1)^+$. Thus, $s^{\mathcal{I}} = \text{Elements}(s_1)^+ = \text{Elements}(s)$.

Case 6 (s is $s_1 \oplus s_2$, where $\oplus \in \{\sqcap, \sqcup, \setminus, \bowtie, *\}$). For all other operators, by the induction hypothesis, we have that $\text{Elements}(s_1) = s_1^{\mathcal{I}}$ and $\text{Elements}(s_2) = s_2^{\mathcal{I}}$. We have that $s^{\mathcal{I}} = s_1^{\mathcal{I}} \oplus s_2^{\mathcal{I}} = \text{Elements}(s_1) \oplus \text{Elements}(s_2)$. For each operator \oplus ,

the upward rule of the calculus for \oplus ensures that $\text{Elements}(s_1) \oplus \text{Elements}(s_2) \subseteq \text{Elements}(s)$ and the downward rule of the calculus for \oplus ensures that $\text{Elements}(s) \subseteq \text{Elements}(s_1) \oplus \text{Elements}(s_2)$. For example, for \sqcup , since UNION UP does not apply to \mathcal{S} , we have that for each $x \in s_i$ in \mathcal{S}^* where $i \in \{1, 2\}$, it must be the case that $x \in s_1 \sqcup s_2$ is in \mathcal{S} as well, and thus $\text{Elements}(s_1) \sqcup \text{Elements}(s_2) \subseteq \text{Elements}(s)$. Since UNION DOWN does not apply to \mathcal{S} , we have that for each $x \in s_1 \sqcup s_2$, it must be the case that $x \in s_i$ for some $i \in \{1, 2\}$, and thus $\text{Elements}(s) \subseteq \text{Elements}(s_1) \sqcup \text{Elements}(s_2)$. Similar reasoning shows $\text{Elements}(s_1) \oplus \text{Elements}(s_2) = \text{Elements}(s)$ for the other operators, and thus $s^{\mathcal{I}} = \text{Elements}(s)$.

Thus, $\text{Elements}(s) = s^{\mathcal{I}}$ for all terms s of set sort in $\mathcal{T}(\mathcal{S})$. We now show that all constraints in \mathcal{S} are satisfied by \mathcal{I} .

(Dis)equalities between non-set variables) Since \mathcal{S}^* is a superset of \mathcal{S} , all equalities between non-set variables $x \approx y$ in \mathcal{S} are satisfied due to the second part of our construction of \mathcal{I} . Since the EQ UNSAT does not apply to \mathcal{S} , all disequalities $x \not\approx y$ between non-set variables in \mathcal{S} are satisfied due to the second part of our construction of \mathcal{I} as well.

(Dis)equalities between tuples) If $\langle x_1, \dots, x_n \rangle \approx \langle y_1, \dots, y_n \rangle \in \mathcal{S}$, then $x_1 \approx y_1, \dots, x_n \approx y_n \in \mathcal{S}^*$ by the first part of our construction of \mathcal{S}^* . By assumption, $x_1, \dots, x_n, y_1, \dots, y_n$ are variables whose sort is not a tuple or set. Thus for the same reasoning as above, we have that $x_1^{\mathcal{I}} = y_1^{\mathcal{I}}, \dots, x_n^{\mathcal{I}} = y_n^{\mathcal{I}}$ and hence $\langle x_1, \dots, x_n \rangle^{\mathcal{I}} = \langle y_1, \dots, y_n \rangle^{\mathcal{I}}$. By assumption, \mathcal{S} contains no disequalities between tuples.

(Dis)equalities between set variables) For each $s_1 \approx s_2 \in \mathcal{S}$ where s_1 and s_2 are set terms, by the third part of our construction of \mathcal{S}^* , we have that $x \in s_1 \in \mathcal{S}^*$ if and only if $x \in s_2 \in \mathcal{S}^*$ for all terms x . Hence, $\text{Elements}(s_1) = \text{Elements}(s_2)$. Since $\text{Elements}(s) = s^{\mathcal{I}}$ for all set terms s in $\mathcal{T}(\mathcal{S})$, we have that $s_1^{\mathcal{I}} = \text{Elements}(s_1) = \text{Elements}(s_2) = s_2^{\mathcal{I}}$, and hence $\mathcal{I} \models s_1 \approx s_2$. To show all disequalities $s_1 \not\approx s_2$ between set terms in \mathcal{S} are satisfied, note that since rule SET DISEQ does not apply to \mathcal{S} , we have $z \in s_1 \in \mathcal{S}$ and $z \notin s_2 \in \mathcal{S}$ for some z , or vice versa. Since rule EQ UNSAT does not apply, we have that $\text{Elements}(s_1) \neq \text{Elements}(s_2)$, $s_1^{\mathcal{I}} \neq s_2^{\mathcal{I}}$ and hence $\mathcal{I} \models s_1 \not\approx s_2$.

Set memberships) For all $x \in s \in \mathcal{S}$, we have that $x \in s \in \mathcal{S}^*$ and thus $x^{\mathcal{I}} \in \text{Elements}(s)$ by construction. Since $\text{Elements}(s) = s^{\mathcal{I}}$ for all terms s of set sort in $\mathcal{T}(\mathcal{S})$, we have that $x^{\mathcal{I}} \in s^{\mathcal{I}}$, and hence $\mathcal{I} \models x \in s$. For all $x \notin s \in \mathcal{S}$, since EQ UNSAT does not apply to \mathcal{S} , we have that $x \in s \notin \mathcal{S}^*$ and hence $x^{\mathcal{I}} \notin \text{Elements}(s)$. Since $\text{Elements}(s) = s^{\mathcal{I}}$ for all terms s of set sort in $\mathcal{T}(\mathcal{S})$, we have that $x^{\mathcal{I}} \notin s^{\mathcal{I}}$, and hence $\mathcal{I} \models x \notin s$.

Thus, \mathcal{I} satisfies all constraints in \mathcal{S} . Since \mathcal{S} is a superset of \mathcal{S}_0 , we have that \mathcal{I} satisfies \mathcal{S}_0 as well. \square

Our calculus is terminating for a sublanguage of constraints involving only unary and binary relations and excluding transitive closure, product, or equality between relations. While this sublanguage, defined in Figure 4, is quite restricted, it is useful in reductions of description logics to relational logic, which we discuss in Section 5.2.

Proposition 3 (Termination). *If \mathcal{S} is a finite set of constraints generated by the grammar in Figure 4, then all derivation trees with root node \mathcal{S} are finite.*

(element)	$e := x$
(unary relation)	$u := x \mid [] \mid u_1 \sqcup u_2 \mid u_1 \sqcap u_2 \mid [\langle e \rangle] \mid b \bowtie u$
(binary relation)	$b := x \mid [] \mid b_1 \sqcup b_2 \mid b_1 \sqcap b_2 \mid [\langle e_1, e_2 \rangle] \mid b^{-1}$
(constraint)	$\varphi := e_1 \approx e_2 \mid \langle e \rangle \in u \mid \langle e_1, e_2 \rangle \in b \mid \neg\varphi_1$

Fig. 4: A restricted fragment of T_R -constraints. Letter x denotes variables.

Proof. Assume that \mathcal{S} is a finite set containing only constraints φ from the grammar in Figure 4. First, we construct the following mapping from relation terms to tuple terms. Let D_u (resp. D_b) be a mapping from unary (resp. binary) relation terms to sets of unary (resp. binary) tuple terms defined as the least solution to the following set of constraints, where the e 's, the u 's and the b 's are implicitly universally quantified metavariables ranging over terms respectively of element, unary relation and binary relation sort:

$$\begin{aligned}
\langle e \rangle \in D_u(u) & \text{ if } \langle e \rangle \in u \in \mathcal{S} \\
\langle e \rangle \in D_u(u_1) & \text{ if } \langle e \rangle \in D_u(u_2) \text{ and } u_1 \in \mathcal{T}(u_2) \\
\langle z_{e,b,u} \rangle \in D_u(u) & \text{ if } \langle e \rangle \in D_u(b \bowtie u) \\
\langle e_1, e_2 \rangle \in D_b(b) & \text{ if } \langle e_1, e_2 \rangle \in b \in \mathcal{S} \\
\langle e_1, e_2 \rangle \in D_b(b_1) & \text{ if } \langle e_i, e_j \rangle \in D_b(b_2) \text{ for } \{i, j\} = \{1, 2\} \text{ and } b_1 \in \mathcal{T}(b_2) \\
\langle e, z_{e,b,u} \rangle \in D_b(b) & \text{ if } \langle e \rangle \in D_u(b \bowtie u)
\end{aligned}$$

where $z_{e,b,u}$ denotes a unique fresh variable for each value of e , b , and u . We require only one such variable for each triple (e, b, u) since our redundancy criteria for rule applications ensures that JOIN DOWN cannot be applied more than once for the same premise $\langle e \rangle \in D_u(b \bowtie u)$. Intuitively, D_u maps each relation term u in \mathcal{S} to an overapproximation of the set of unary tuples $\langle e \rangle$ for which our calculus can infer the constraint $\langle e \rangle \in u$ using downward rules only, and similarly for the binary case D_b . The domain of D_u and that of D_b contain only relation terms occurring in \mathcal{S} , and thus are finite. All sets in the ranges of D_u and D_b are also finite. To show this, we argue that only a finite number of fresh variables $z_{e,b,u}$ are introduced by this construction. We define a measure depth on element terms such that $\text{depth}(e) = 0$ for all $e \in \mathcal{T}(\mathcal{S})$, and $\text{depth}(z_{e,b,u}) = 1 + \text{depth}(e)$. For all variables $z_{e,b,u}$ in the range of D_u and D_b , we have that $b \bowtie u \in \mathcal{T}(\mathcal{S})$, and if e is a variable of the form $z_{e',b',u'}$, then $b \bowtie u$ is either a subterm of b' or u' . Thus, the depth of all element terms in the range of D_u and D_b is finite. Since there are finitely many element terms in $\mathcal{T}(\mathcal{S})$, and finitely many terms of the form $b \bowtie u$ in $\mathcal{T}(\mathcal{S})$, there are finitely many variables of the form $z_{e,b,u}$ and thus finitely many element terms occur in tuple terms in the range of D_u and D_b . Therefore, there are finitely many tuple terms in sets in the ranges of D_u and D_b .

Now, let U_u (resp. U_b) be a mapping from unary (resp. binary) relation terms to sets of unary (resp. binary) tuple terms, constructed to be the least solution to the following set of constraints (where again the e 's, the u 's and the b 's are implicitly universally

quantified metavariables as above):

$$\begin{aligned}
\langle e \rangle \in U_u(u) & \text{ if } \langle e \rangle \in D_u(u) \\
\langle e \rangle \in U_u(u_1) & \text{ if } \langle e \rangle \in U_u(u_2), u_2 \in \mathcal{T}(u_1) \text{ and } u_1 \in \mathcal{T}(\mathcal{S}) \\
\langle e \rangle \in U_u(u) & \text{ if } u = [e] \text{ and } u \in \mathcal{T}(\mathcal{S}) \\
\langle e_1 \rangle \in U_u(u_1) & \text{ if } u_1 = b_1 \bowtie u_2 \text{ and } \langle e_1, e_2 \rangle \in U_b(b_1) \\
\langle e_1, e_2 \rangle \in U_b(b) & \text{ if } \langle e_1, e_2 \rangle \in D_b(b) \\
\langle e_1, e_2 \rangle \in U_b(b_1) & \text{ if } \langle e_i, e_j \rangle \in U_b(b_2) \text{ for } \{i, j\} = \{1, 2\}, b_2 \in \mathcal{T}(b_1), b_1 \in \mathcal{T}(\mathcal{S}) \\
\langle e_1, e_2 \rangle \in U_b(b) & \text{ if } b = [\langle e_1, e_2 \rangle] \text{ and } b \in \mathcal{T}(\mathcal{S})
\end{aligned}$$

Similar to the previous construction, U_u maps each relation term u in \mathcal{S} to an over-approximation of the set of unary tuples $\langle e \rangle$ for which our calculus can infer the constraint $\langle e \rangle \sqsubseteq u$ using both downward and upward rules, and similarly for the binary case U_b . By construction, since the domains of D_u and D_b are subsets of $\mathcal{T}(\mathcal{S})$, the domains of U_u and U_b are also subsets of $\mathcal{T}(\mathcal{S})$, and thus are finite, hence their respective ranges R_u and R_b are finite too. Each set in R_u or R_b is finite as well, since the tuples in these ranges are built from element terms e that occur in the range of D_u and D_b or in singleton sets of the form $[\langle e \rangle]$, $[\langle e, e' \rangle]$ or $[\langle e', e \rangle]$ in $\mathcal{T}(\mathcal{S})$.

Now let $\widehat{\mathcal{S}}$ be the following set of constraints:

$$\begin{aligned}
\widehat{\mathcal{S}} = & \{(\neg)\langle e \rangle \sqsubseteq u \mid \langle e \rangle \in U_u(u), u \in \mathcal{T}(\mathcal{S})\} \cup \\
& \{(\neg)\langle e_1, e_2 \rangle \sqsubseteq b \mid \langle e_1, e_2 \rangle \in U_b(b), b \in \mathcal{T}(\mathcal{S})\} \cup \\
& \{(\neg)e_1 \approx e_2 \mid e_1, e_2 \in \mathcal{T}(R_u \cup R_b)\} \cup \\
& \{\langle e_1 \rangle \approx \langle e_2 \rangle \mid e_1, e_2 \in \mathcal{T}(R_u \cup R_b)\} \cup \\
& \{\langle e_1, e_2 \rangle \approx \langle e_3, e_4 \rangle \mid e_1, e_2, e_3, e_4 \in \mathcal{T}(R_u \cup R_b)\}
\end{aligned}$$

From the arguments above we can conclude that $\widehat{\mathcal{S}}$ is finite. By construction, $\mathcal{S} \subseteq \widehat{\mathcal{S}}$. One can show by structural induction on derivation trees that all descendants of \mathcal{S} in a derivation tree are also subsets of $\widehat{\mathcal{S}}$. Since the size of a state strictly grows with each rule application not deriving *unsat*, we can conclude that no derivation tree can be grown indefinitely. Hence all derivation trees with root \mathcal{S} are finite. \square

By Proposition 1, 2, and 3, we have that *any* rule application strategy for the calculus is a decision procedure for finite sets of constraints in the language generated by the grammar from Figure 4.

5 Applications of T_R

Our main motivation for adding native support for relations in an SMT solver is that it enables more natural mappings from other logical formalisms ultimately based on relations. This opens the possibility of leveraging the power and flexibility of SMT to reason about problems expressed in those formalisms. We discuss here two potential applications: reasoning about Alloy specifications and reasoning about OWL ontologies. It should be clear to the knowledgeable reader though that the set of potential applications is much larger, encompassing description logics in general as well as various modal logics—via an encoding of their accessibility relation.

5.1 Alloy specifications

Alloy is a formal specification language based on relational logic which is widely used for modeling structurally-rich problems [12]. Alloy specifications, called *models* in the Alloy literature, are built from relations and relational algebra operations in addition to the usual logical connectives and quantifiers. One can also specify expected properties of a specification as formulas, called *assertions* in Alloy, that should be entailed by the specification.

The analysis of Alloy specifications can be performed automatically by a tool called the Alloy Analyzer which uses as its reasoning engine Kodkod, a SAT-based finite model finder [17]. This requires the user to impose a (concrete, artificial) finite upper bound on the size of the domains of each relation, limiting the analyzer’s ability to determine that a specification is consistent or has a given property. In the first case, the user has to manually increase the bounds until the Alloy Analyzer is able to find a satisfying interpretation for the specification; in the second case, until it can find a counter-example for the property. As a consequence, the analyzer cannot be used to prove that (i) a specification is inconsistent or (ii) it does have a certain property.

In contrast, thanks to its new theory solver for relational constraints based on the calculus described earlier, CVC4 is now able in many cases to do (i) and (ii) automatically, with no artificial upper bounds on domain sizes. Also, because of its own finite model finding capabilities [15], it can find minimal satisfying interpretations for consistent specifications or minimal counter-examples for properties without the need of user-provided artificial upper bounds on domain sizes. Finally, since its relational theory solver is fully integrated with its theory solvers for other theories (such as linear arithmetic, strings, arrays, and so on), CVC4 can natively support mixed constraints using relations over its various built-in types, something that is possible in the Alloy Analyzer only in rather limited form.⁵ To evaluate CVC4’s capabilities in solving Alloy problems, we have defined a translation from Alloy specifications to semantically equivalent SMT formulas that leverages our theory of relations. The translation focuses on Alloy’s kernel language since non-kernel features can be rewritten to the kernel language by the Alloy Analyzer itself. We sketch our translation below.⁶

In Alloy, a *signature* is a set of uninterpreted atomic elements, called *atoms*. Signatures are defined with a syntax that is reminiscent of classes in object-oriented languages. A *relation* (of arity n) is a set of n -tuples of atoms and is declared as a *field* of some signature S , which acts as the domain of the elements in the first component. Multiplicity constraints on signatures and fields can be added with keywords such as **some**, **no**, **lone**, and **one**, which specify that a signature is non-empty, empty, has cardinality at most one, and is a singleton, respectively. Other keywords specify that a relation is one-to-one, one-to-many, and so on. One or more signatures can be declared to be subsets of another signature with **extends** or **in**. With **extends**, all specified signatures are additionally assumed to be mutually disjoint.

⁵ The Alloy Analyzer currently has built-in support for bounded integers. Any other data types need to be axiomatized in the specification.

⁶ The translation is sound only if all Alloy signatures are assumed to be finite. A full account of the translation and a proof of its soundness are beyond the scope of this paper.

In the translation, we introduce an uninterpreted sort `Atom` for Alloy atoms. An Alloy signature $\mathbf{sig} \mathbf{s}$ is translated as a constant⁷ S of sort $\text{Rel}_1(\text{Atom})$, that is, a set of unary tuples. A field $\mathbf{f}: \mathbf{s}_1 \rightarrow \dots \rightarrow \mathbf{s}_n$ of a signature \mathbf{s} is translated as a constant f of sort $\text{Rel}_{n+1}(\text{Atom}, \dots, \text{Atom})$ together with the additional constraint $f \sqsubseteq S * S_1 * \dots * S_n$ to ensure that the components of f 's tuples are from the intended signatures. The signature hierarchy is encoded using subset constraints. For example, the Alloy constraint $\mathbf{sig} \mathbf{s}_1, \dots, \mathbf{s}_n \text{ extends } \mathbf{s}$ is translated as the set of constraints $\{S_1 \sqsubseteq S, \dots, S_n \sqsubseteq S\} \cup \{S_i \sqcap S_j \approx [] \mid 1 \leq i < j \leq n\}$. If \mathbf{s} above is also declared to be *abstract* (a notion similar to abstract classes in object-oriented languages), the additional constraint $S_1 \sqcup \dots \sqcup S_n \approx S$ is added to enforce that. Similarly, signature declarations of the form $\mathbf{sig} \mathbf{s}_1, \dots, \mathbf{s}_n \text{ in } \mathbf{s}$, are translated just as $\{S_1 \sqsubseteq S, \dots, S_n \sqsubseteq S\}$. Multiplicity constraints are translated as quantified formulas.

Since our theory supports all constructs and operators in Alloy's kernel language, Alloy expressions and formulas, which can include quantifiers ranging over atoms, can be more or less directly translated to their counterparts in CVC4's language. It is worth mentioning that our translation supports Alloy's set comprehension construct, by introducing a fresh relational constant for the set and adding definitional axioms for it. In addition, we partially support Alloy cardinality constraints of the form $\#(r) \text{ op } n$ where r is a relation term, $\text{op} \in \{<, >, =\}$, $n \in \mathbb{N}$, and $\#$ is the cardinality operator, by encoding them as subset constraints. For example, the Alloy constraint $\#(\mathbf{s}) < 3$ on a signature S is translated to $S \sqsubseteq [\langle k_1 \rangle, \langle k_2 \rangle]$ where S is the corresponding unary relation and k_1 and k_2 are two fresh constants of sort `Atom`.

5.2 OWL DL Ontologies

OWL is an ontology language whose current version, OWL 2, was adopted as a standard Semantic Web language by the W3 consortium. It includes a sublanguage, called OWL DL, that corresponds to the expressive, yet decidable, description logic $\mathcal{SHOIN}(\mathbf{D})$ [2]. We have defined an initial, partial translation from $\mathcal{SHOIN}(\mathbf{D})$ to SMT formulas that again leverages our theory of relations.

A mapping from salient $\mathcal{SHOIN}(\mathbf{D})$ constructs to their SMT counterparts is illustrated in Figure 5. The figure shows only relations whose elements do not belong to the so-called *concrete domain(s)* \mathbf{D} of $\mathcal{SHOIN}(\mathbf{D})$.⁸ As with the Alloy translation, we use the single sort `Atom` for all elements of non-concrete domains. The set comprehension notation is used here for brevity: for a set comprehension term of the form $[x \mid \varphi]$ where x has n -tuple sort, we introduce a fresh set constant S accompanied by the defining axiom $\forall x: \text{Tuple}_n(\text{Atom}) (x \in S \Leftrightarrow \varphi)$. The constant `Univ` in the figure denotes the universal unary relation over `Atom`. Since this constant is currently not built-in as a symbol of T_R , it is accompanied by the defining axiom $\forall a: \text{Atom} \langle a \rangle \in \text{Univ}$. The expression $\text{dist}(a_1, \dots, a_n)$ states that a_1, \dots, a_n are pairwise different. We observe how the translation is immediate for most constructs, with the notable exception of universal and number restrictions, which require the use of complex quantified formulas.

⁷ Free constants have the same effect as free variables for satisfiability purposes.

⁸ Some of those domains in OWL correspond to built-in sorts in CVC4. A full translation from OWL concrete domains to CVC4 built-in sorts is beyond the scope of this work.

DL	CVC4
individual name: a	$a : \text{Atom}$
atomic concept C, role R	$C : \text{Rel}_1(\text{Atom}), R : \text{Rel}_2(\text{Atom}, \text{Atom})$
intersection $C \sqcap D$, union $C \sqcup D$	$C \sqcap D, C \sqcup D$
inverse role R^- , complement $\neg C$	$R^{-1}, \text{Univ} \setminus C$
top concept \top , bottom concept \perp	$\text{Univ}, []$
existential restriction $\exists R.C$	$R \bowtie C$
universal restriction $\forall R.C$	$[x \mid x \in \text{Univ} \wedge [x] \bowtie R \sqsubseteq C]$
at-least restriction $\geq_n R.C$	$[x \mid x \in \text{Univ} \wedge \exists a_1, \dots, a_n : \text{Atom}([x] \bowtie R) \sqcap C \sqsupseteq [\langle a_1 \rangle, \dots, \langle a_n \rangle] \wedge \text{dist}(a_1, \dots, a_n)]$
at-most restriction $\leq_n R.C$	$[x \mid x \in \text{Univ} \wedge \exists a_1, \dots, a_n : \text{Atom}([x] \bowtie R) \sqcap C \sqsubseteq [\langle a_1 \rangle, \dots, \langle a_n \rangle] \wedge [\langle a_1 \rangle, \dots, \langle a_n \rangle] \sqsubseteq C]$
local reflexivity $\exists R.\text{Self}$	$[\langle x, y \rangle \mid \langle x, y \rangle \in R \wedge x \approx y]$
nominal $\{a\}$	$[\langle a \rangle]$
concept, role assertion $C(a), R(a, b)$	$a \in C, \langle a, b \rangle \in R$
individual (dis)equality $a \approx b, a \not\approx b$	$a \approx b, a \not\approx b$
concept, role inclusion $C \sqsubseteq D, R \sqsubseteq S$	$C \sqsubseteq D, R \sqsubseteq S$
concept, role equiv. $C \equiv D, R \equiv S$	$C \approx D, R \approx S$
complex role inclusion $R_1 \circ R_2 \sqsubseteq S$	$R_1 \bowtie R_2 \sqsubseteq S$
role disjointness $\text{Disjoint}(R, S)$	$R \sqcap S \approx []$

Fig. 5: A mapping from DL language to Σ_R -constraints.

6 Evaluation

To evaluate our theory solver for T_R in CVC4 we implemented translators from Alloy and from OWL based on the translations sketched in the previous section. This section presents an initial evaluation on a selection of Alloy and OWL benchmarks.⁹

6.1 Experimental Evaluation on Alloy Models

We considered two sets of Alloy benchmarks; the first consists of 40 examples from the Alloy distribution and from a formal methods course taught by one of the authors; the second were used in [10] to evaluate AlloyPE. All benchmarks consist of an Alloy model together with a single property. We evaluated two configurations of CVC4. The first, denoted **CVC4**, enables full native support for relational operators via the calculus from Section 3. The second, denoted **CVC4+AX**, instead encodes all relational operators as uninterpreted functions and supplements the translation of benchmarks with additional axioms that specify their semantics. To compare CVC4 with other tools, we also evaluated the Alloy Analyzer, version 4.2, downloaded from the Alloy website, and

⁹ Detailed results and all benchmarks are available at <http://cvc4.cs.stanford.edu/papers/CADE2017-relations/>.

Problem	Alloy Analyzer		CVC4		CVC4+AX		AlloyPE	
	Res.	Time/Scope	Res.	Time	Res.	Time	Res.	Time
academia_0	sat	0.60/3	sat	1.55	-	to	unk	84.76
academia_1	sat	0.53/2	sat	1.93	-	to	-	to
academia_2	sat	0.45/2	sat	0.49	-	to	unk	0.15
social_1	sat	0.52/3	sat	1.20	-	to	n/a	-
social_5	sat	1.56/2	sat	0.49	-	to	n/a	-
social_6	sat	0.49/2	sat	0.52	-	to	n/a	-
cf_0	sat	0.47/3	sat	0.51	-	to	n/a	-
cf_1	sat	0.49/3	sat	0.78	-	to	n/a	-
javatypes	sat	0.50/3	sat	0.42	-	to	uns	2.35
set	sat	0.45/2	sat	0.46	-	to	unk	0.92
loc_int	sat	0.57/1	sat	2.82	-	to	n/a	-
genealogy	sat	0.64/6	sat	89.20	-	to	n/a	-
number_1	sat	0.81/2	sat	8.65	-	to	n/a	-
railway	sat	0.67/4	sat	156.45	-	to	n/a	-
academia_3	b-uns	162.17/63	uns	0.49	uns	1.05	uns	0.28
academia_4	b-uns	246.92/162	uns	0.43	uns	0.54	uns	0.13
family_1	b-uns	146.62/68	uns	0.41	uns	0.44	uns	0.15
family_2	b-uns	279.77/30	uns	1.02	uns	48.78	uns	0.23
social_2	b-uns	256.98/56	uns	0.66	-	to	n/a	-
social_3	b-uns	191.45/57	uns	0.49	uns	35.91	n/a	-
social_4	b-uns	171.26/64	uns	0.46	uns	18.13	n/a	-
birthday	b-uns	156.08/53	uns	0.45	uns	0.61	uns	0.13
library	b-uns	259.54/119	uns	0.42	uns	0.40	uns	1.11
lights	b-uns	228.89/122	uns	32.69	-	to	n/a	-
INSLabel	b-uns	198.53/8	uns	1.46	-	to	n/a	-
farmers_1	sat	1.04/8	-	to	-	to	n/a	-
views	sat	9.91/9	-	to	-	to	n/a	-

Fig. 6: Evaluation on Alloy benchmarks.

El Ghazi et al.’s AlloyPE tool (kindly provided to us directly by its authors) using the SMT solver Z3 version 4.5.1 as a backend. All experiments were performed with a 300 second timeout on a machine with a 2.9 GHz Intel Core i7 CPU with 8 GB of memory.

Figures 6 and 7 show the results from running the Alloy Analyzer, CVC4 and AlloyPE on the two sets of Alloy benchmarks. We omit results for 13 of the benchmarks from the first set that no system solved. The second and third columns show the results of running the Alloy Analyzer. To evaluate the Alloy Analyzer on these benchmarks, we considered bounded scopes in an incremental fashion. Using a script, we set an initial upper bound, or *scope*, of 1 for the cardinality of all signatures in the problem, and kept increasing it by 1 if the Alloy Analyzer found the problem unsatisfiable (b-uns) in the current scope—meaning that it was not able to *disprove* the property in the problem. We terminated on time out (to), or when the analyzer was able to disprove the property (sat) or ran out of memory for a scope. We report the scope size for benchmarks where the tool returned an answer. The fourth and fifth columns show the results from CVC4 when invoked in finite model finding mode on the translated SMT problem. The last two columns are the results from AlloyPE, where n/a indicates that AlloyPE failed due either to the presence of unsupported Alloy constructs in the input problem or to internal errors during solving or translation.

As shown in the table, our approach is overall slower than the Alloy Analyzer for satisfiable benchmarks. For the unsatisfiable ones, CVC4 returns an answer within a reasonable time limit for most of the benchmarks and has advantages over the state of the art. It is important to note, however, that an unsat answer from CVC4 indicates that

Problem	Alloy Analyzer		CVC4		CVC4+AX		AlloyPE	
	Res.	Time	Res.	Time	Res.	Time	Res.	Time
mem-wr	b-uns	195.98/35	uns	0.43	uns	0.48	uns	0.44
mem-wi	b-uns	260.66/29	uns	0.45	uns	0.50	uns	0.42
ab-ai	b-uns	185.06/28	uns	0.46	uns	0.79	uns	0.49
ab-dua	b-uns	193.33/27	uns	0.49	uns	0.48	uns	0.70
abt-dua	b-uns	137.87/14	uns	0.60	uns	0.81	uns	0.70
abt-ly-u	b-uns	261.23/9	uns	0.81	uns	28.26	uns	1.4
abt-ly-p	b-uns	277.86/8	uns	0.81	uns	1.77	uns	175.19
gp-nsf	b-uns	152.55/69	uns	0.41	uns	0.59	uns	0.43
gp-nsq	b-uns	166.75/66	uns	0.42	-	to	uns	0.44
com-1	b-uns	297.18/13	uns	2.95	-	to	uns	0.59
com-2	b-uns	295.73/13	-	to	-	to	uns	0.55
com-3	b-uns	295.33/14	uns	4.29	-	to	uns	0.64
com-4a	b-uns	301.57/13	uns	9.39	-	to	uns	0.99
com-4b	b-uns	299.77/13	uns	0.90	-	to	uns	0.61
fs-sd	b-uns	157.90/70	uns	0.42	-	to	uns	0.89
fs-nda	b-uns	271.38/44	uns	0.55	-	to	uns	0.83
gc-s1	b-uns	270.07/14	uns	4.92	uns	8.14	uns	14.27
gc-s2	b-uns	288.44/8	-	to	-	to	uns	10.66
gc-c	b-uns	287.73/8	-	to	-	to	uns	42.31
hr-l	b-uns	275.80/7	-	to	-	to	-	to

Fig. 7: Evaluation on AlloyPE benchmarks

the property is valid as opposed to the b-uns answer from the Alloy Analyzer, which only means the property is valid within the given scope.

Compared to AlloyPE, we successfully solved all of their benchmarks but four, as indicated in Figure 7. For these benchmarks, AlloyPE benefits from performing a static analysis of the problem that involves sophisticated heuristics to discover invariants. For our own set of benchmarks, in Figure 6, AlloyPE failed on all sat benchmarks and was unable to solve many unsat ones due to failures during the translation or the solving phase. We observe that AlloyPE gives an unsound answer for the benchmark javatypes: it returns unsat, whereas both the Alloy Analyzer and CVC4 return sat.

Our results also indicate that native support for relational reasoning is important for reasoning efficiently for these benchmark sets. In fact, **CVC4+AX** is unable to report sat for any satisfiable benchmark due to its use of axioms for the relation operators, which quantify over set variables, where CVC4’s finite model finding techniques are not applicable. More interestingly, **CVC4+AX** solves significantly fewer unsat benchmarks when compared to **CVC4**, indicating that using the calculus in Section 3 is superior to encoding the semantics of relational operators via an explicit axiomatization.

6.2 Experimental Evaluation on OWL Models

We built a preliminary translator from OWL to SMT, and we did a *consistency evaluation*, which checks whether or not an ontology is contradictory, for a set of OWL benchmarks in pure description logic from the 4th OWL Reasoner Evaluation competition.¹⁰ Of the original 7,704 benchmarks, we considered only those whose size was under 1 MB, and further excluded benchmarks that involved some of the more sophisticated features of the OWL language that are currently not supported by our translator.

¹⁰ See <https://www.w3.org/community/owled/ore-2015-workshop/competition>.

We ran the experiments with a 30 second time out, on a machine with a 3.2GHz Intel(R) Xeon CPU E5-2667 v3 and 20 GB of memory.

Among the selected 3,936 benchmarks, CVC4 found 3,639 consistent, found 7 inconsistent, and timed out on the remaining 290. By comparison, the state-of-the-art DL reasoner Konclude [16] gave an answer for all 3,936 benchmarks. However, Konclude and CVC4 disagreed on 9 benchmarks, all of which CVC4 found consistent. We determined that Konclude reports inconsistent for those 9 benchmarks possibly because the benchmarks are not syntactically compliant. For example, some include Boolean literals without a type declaration. In terms of performance, CVC4 takes on average 1.7 seconds per benchmark on the 3,646 it solves. This is significantly slower than Konclude, which takes on average 0.02 seconds per benchmark on the same set. We attribute this to the fact that CVC4 does not yet support the universal set and set complement natively, and has no specific quantifier instantiation heuristics for the quantified formulas generated by the translation of universal and number restrictions. Nevertheless, we find these results quite encouraging as they show that further investigation into efficient reasoning for OWL models in SMT solvers is an interesting direction of research.

7 Conclusion and Future Work

We presented a calculus for an extension to the theory of finite sets that includes support for relations and relational operators. We implemented the calculus as a modular extension to the set subsolver in our SMT solver CVC4. A preliminary evaluation has shown that our implementation is competitive with the state of the art when used to prove properties or verify the consistency of Alloy specifications.

We are investigating more expressive fragments for which our calculus terminates, including those corresponding to fragments of description logic [11]. In future work, we would like to devise an approach for a theory that includes both relational constraints and cardinality constraints that is efficient in practice, together with specialized techniques geared toward reasoning about formulas resulting from the translation of description logic problems. In particular, we plan to extend our logic with the set complement operator and a constant for the universal set, and extend our calculus and its implementation to provide direct support for them.

Acknowledgements This work was partially supported by NSF grant no. 1228765 and by a gift from GE Global Research. We are grateful to Jasmin Blanchette and the anonymous reviewers for their very detailed comments and questions which helped improve the presentation of the paper.

References

1. F. Baader. *The description logic handbook: Theory, implementation and applications*. Cambridge university press, 2003.
2. F. Baader, I. Horrocks, and U. Sattler. Description logics. In V. L. Frank van Harmelen and B. Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 135 – 179. Elsevier, 2008.

3. K. Bansal, A. Reynolds, C. W. Barrett, and C. Tinelli. A new decision procedure for finite sets and cardinality constraints in SMT. In *Proceedings of IJCAR'16*, volume 9706 of *LNCS*, pages 82–98. Springer, 2016.
4. C. Barrett, C. L. Conway, M. Deters, L. Hadarean, D. Jovanović, T. King, A. Reynolds, and C. Tinelli. CVC4. In *Proceedings of CAV'11*, volume 6806 of *LNCS*, pages 171–177. Springer, 2011.
5. C. Barrett, P. Fontaine, and C. Tinelli. The SMT-LIB standard—Version 2.6. In A. Gupta and D. Kroening, editors, *SMT 2010*, 2010.
6. C. Barrett, R. Sebastiani, S. Seshia, and C. Tinelli. Satisfiability modulo theories. In A. Biere, M. J. H. Heule, H. van Maaren, and T. Walsh, editors, *Handbook of Satisfiability*, volume 185, chapter 26, pages 825–885. IOS Press, February 2009.
7. B. Dutertre and L. D. Moura. The YICES SMT solver. Technical report, SRI International, 2006.
8. A. A. E. Ghazi and M. Taghdiri. Analyzing alloy constraints using an SMT solver: a case study. In *5th International Workshop on Automated Formal Methods (AFM)*, 2010.
9. A. A. E. Ghazi and M. Taghdiri. Relational reasoning via SMT solving. In *Proceedings of FM'11*, volume 6664 of *LNCS*, pages 133–148. Springer, 2011.
10. A. A. E. Ghazi, M. Taghdiri, and M. Herda. First-order transitive closure axiomatization via iterative invariant injections. In *Proceedings of NFM'15*, volume 9058 of *LNCS*. Springer, 2015.
11. I. Horrocks and U. Sattler. Decidability of shiq with complex role inclusion axioms. *Artificial Intelligence*, 160(1-2):79–104, 2004.
12. D. Jackson. Alloy: a lightweight object modelling notation. *ACM Trans. Softw. Eng. Methodol.*, 11(2):256–290, 2002.
13. D. Jackson. *Software Abstractions - Logic, Language, and Analysis*. MIT Press, 2006.
14. R. Nieuwenhuis, A. Oliveras, and C. Tinelli. Solving SAT and SAT Modulo Theories: from an abstract Davis-Putnam-Logemann-Loveland Procedure to DPLL(T). *Journal of the ACM*, 53(6):937–977, Nov. 2006.
15. A. Reynolds, C. Tinelli, A. Goel, and S. Krstic. Finite model finding in SMT. In *Proceedings of CAV'13*, volume 8044 of *LNCS*, pages 640–655. Springer, 2013.
16. A. Steigmiller, T. Liebig, and B. Glimm. Konclude: System description. *Web Semantics: Science, Services and Agents on the World Wide Web*, 27(1), 2014.
17. E. Torlak and D. Jackson. Kodkod: a relational model finder. In *Proceedings of TACAS'07*, volume 4424 of *LNCS*, pages 632–647. Springer, 2007.
18. D. Tsarkov and I. Horrocks. Fact++ description logic reasoner: system description. In *Proceedings of IJCAR'06*, volume 4130 of *LNCS*. Springer, 2006.
19. D. Tsarkov and I. Palmisano. Chainsaw: a metareasoner for large ontologies. In I. Horrocks, M. Yatskevich, and E. Jiménez-Ruiz, editors, *ORE*, 2012.
20. W3C. OWL 2 web ontology language, <https://www.w3.org/2007/OWL/wiki/Syntax>.