

E-Matching with Free Variables

Philipp Rümmer

Department of Information Technology, Uppsala University, Sweden

Abstract. E-matching is the most commonly used technique to handle quantifiers in SMT solvers. It works by identifying characteristic sub-expressions of quantified formulae, named triggers, which are matched during proof search on ground terms to discover relevant instantiations of the quantified formula. E-matching has proven to be an efficient and practical approach to handle quantifiers, in particular because triggers can be provided by the user to guide proof search; however, as it is heuristic in nature, e-matching alone is typically insufficient to establish a complete proof procedure. In contrast, free variable methods in tableau-like calculi are more robust and give rise to complete procedures, e.g., for first-order logic, but are not comparable to e-matching in terms of scalability. This paper discusses how e-matching can be combined with free variable approaches, leading to calculi that enjoy similar completeness properties as pure free variable procedures, but in which it is still possible for a user to provide domain-specific triggers to improve performance.

1 Introduction

SAT and SMT solvers form the backbone of many of today’s verification systems, responsible for discharging verification conditions that encode correctness properties of hardware or software designs. Such verification conditions are often generated in the context of intricate theories, including various kinds of arithmetic, uninterpreted functions and equality, the theory of arrays, or the theory of quantifiers. Despite much research over the past years, efficient and scalable reasoning in the combination of such theories remains challenging: in particular for handling quantifiers, most state-of-the-art SMT solvers have to resort to heuristic techniques like e-matching and triggers [7, 8]. E-matching is a popular method due to its simplicity and performance, but offers little completeness guarantees and is sensitive to syntactic manipulations of input formulae.

This paper takes the standpoint that heuristics like e-matching should be considered as *optimisations*, and triggers as *hints*, possibly affecting the performance, but not the completeness of an SMT solver. In other words, the set of formulae that a solver can prove should be independent from chosen triggers. Working towards this goal, the paper presents calculi integrating constraint-based free variable reasoning with e-matching, the individual contributions being (i) a free variable sequent calculus for first-order logic (Sect. 3), with support for e-matching and user-provided triggers to guide instantiation of quantified formulae, partly inspired by the positive unit hyper-resolution calculus [14, 15]; (ii) a

similar calculus for first-order logic modulo linear integer arithmetic (Sect. 5), extending the calculus in [23]; (iii) as a component of both calculi, an approach to capture functions and congruence closure procedures (commonly used in SMT) as uninterpreted predicates (Sect. 4); (iv) a complete implementation of the calculus in (ii), called PRINCESS, and experimental evaluation against SMT solvers competing in the SMT competition 2011 (AUFLIA category) (Sect. 6).

The calculus in (i) is sound and complete for first-order logic, while (ii) is sound and complete for fragments such as Presburger arithmetic, the universal and the existential fragment of first-order logic modulo integers, and the languages accepted by related methods like $\mathcal{ME}(\text{LIA})$ [4] and the complete instantiation method in [9]. The completeness results are significantly stronger than those guaranteed by most SMT solvers, and hold *independently* from the application of e-matching or the choice of triggers in proofs.

1.1 Introductory Example

We start by illustrating e-matching and free variable methods using an example. The first-order theory of non-extensional arrays [16] is often encoded using uninterpreted function symbols sel and sto by means of the following axioms:

$$\forall x, y, z. \underline{sel(sto(x, y, z), y)} \doteq z \quad (1)$$

$$\forall x, y_1, y_2, z. (y_1 \doteq y_2 \vee \underline{sel(sto(x, y_1, z), y_2)}) \doteq sel(x, y_2) \quad (2)$$

Intuitively, $sel(x, y)$ retrieves the element of array x stored at position y , while $sto(x, y, z)$ denotes the array that is identical to x , except that position y stores value z . In order to prove that some formula holds over the theory of arrays, the underlined expressions can be used as *triggers* that determine when and how the axioms should be instantiated. Generally, triggers consist of a single or multiple expressions (normally sub-expressions in the body of the quantified formula) that contain all quantified variables. For instance, to prove that the implication

$$b \doteq sto(a, 1, 2) \rightarrow sel(b, 2) \doteq sel(a, 2) \quad (3)$$

holds over the theory of arrays, we can observe that the term $sel(sto(a, 1, 2), 2)$ occurs in the implication, modulo some equational reasoning. This term matches the underlined pattern in (2), and suggests to instantiate (2) to obtain the instance $1 \doteq 2 \vee \underline{sel(sto(a, 1, 2), 2)} \doteq sel(a, 2)$. In fact, (3) follows for this instance of (2), when reasoning in the theories of uninterpreted functions and arithmetic, which allows us to conclude the validity of (3).

Axioms and triggers as shown above are commonly used in SMT solvers, and give rise to efficient decision procedures for ground problems over arrays.¹ However, in the presence of quantifiers, e-matching might be unable to determine the right instantiations, possibly because required instantiations do not exist as ground terms in the formula. For instance, variants of (3) might include:

$$b \doteq sto(a, 1, 2) \rightarrow \exists x. sel(b, x) \doteq sel(a, 2) \quad (4)$$

$$b \doteq sto(a, 1, 2) \rightarrow \exists x. sel(b, x + 1) \doteq sel(a, 2) \quad (5)$$

$$b \doteq sto(a, 1, 2) \rightarrow \exists x. sel(b, x) \doteq sel(a, x) \quad (6)$$

Although the formulae are still valid, the match $sel(sto(a, 1, 2), 2)$ used previously has been eliminated, which makes proof search more intricate. The state-of-the-art e-matching-based SMT solver CVC3 ([3], version 2.4.1) is able to solve (3), but none of (4), (5), (6). A more realistic example, though similar in nature to the formulae shown here, was reported in [12, Sect. 3.3], where a simple modification (Skolemisation) of a small formula prevented Z3 [19] from finding a proof. The goal of the calculus developed in this paper (and of our implementation PRINCESS) is to obtain a system that is more robust against such modifications, by combining e-matching with constraint-based free variable reasoning, while retaining the scalability of SMT solvers.

The general philosophy of free variable methods [11] is to delay the choice of instantiations for quantified formulae with the help of symbolic reasoning. For example, we could instantiate the formula $\exists x.sel(b, x + 1) \doteq sel(a, 2)$ using a free variable X , resulting in $sel(b, X + 1) \doteq sel(a, 2)$. Modulo equational reasoning, this creates the term $sel(sto(a, 1, 2), X + 1)$, which can be unified with the trigger in (2) under the constraint $X \doteq 1$. It is then possible to proceed with the proof as described above. After closing the proof, we can conclude that (5) indeed holds, since the derived constraint $X \doteq 1$ is satisfiable: it is possible (retrospectively) to instantiate $\exists x.sel(b, x + 1) \doteq sel(a, 2)$ with the concrete term $X = 1$.

This example demonstrates that a free variable calculus can be used to compute answers to queries, in a manner similar to constraint logic programming. The system developed in this paper is more general than “ordinary” logic programming, however, since no restrictions on the use of quantifiers are imposed.

2 Background

2.1 Syntax and Semantics of Considered Logics

We assume familiarity with classical first-order logic (FOL, e.g., [11]). Let x range over an infinite set X of variables, c over an infinite set C of constant symbols, p over a set P of uninterpreted predicates with fixed arity, f over a set F of uninterpreted functions with fixed arity, and α over the set \mathbb{Z} of integers. The syntax of the untyped logics in this paper is defined by the following grammar:

$$\begin{aligned} \phi & ::= \phi \wedge \phi \mid \phi \vee \phi \mid \neg \phi \mid \forall x.\phi \mid \exists x.\phi \mid t \doteq 0 \mid t \leq 0 \mid p(t, \dots, t) \\ t & ::= \alpha \mid c \mid x \mid \alpha t + \dots + \alpha t \mid f(t, \dots, t) \end{aligned}$$

The symbol t denotes terms constructed using functions and arithmetic operations. A formula ϕ is called *closed* if all variables in ϕ are bound by quantifiers, and *ground* if it does not contain variables or quantifiers. A location within a formula ϕ is called *positive* if it is underneath an even number of negations \neg , otherwise *negative*. Simultaneous substitution of terms $\bar{t} = (t_1, \dots, t_n)$ for variables $\bar{x} = (x_1, \dots, x_n)$ in ϕ is denoted by $[\bar{x}/\bar{t}]\phi$; we assume that variable capture

¹ We are grateful to the anonymous referees pointing out that a further trigger (not shown here) is needed in (2) for a complete array procedure.

is avoided by renaming bound variables as necessary. For simplicity, we sometimes write $s \doteq t$ as a shorthand of $1 \cdot s + (-1) \cdot t \doteq 0$. The abbreviation *true* (*false*) stands for $0 \doteq 0$ ($1 \doteq 0$), and implication $\phi \rightarrow \psi$ for $\neg\phi \vee \psi$.

We consider fragments of the syntax shown above, including function-free first-order logic (Sect. 2.3, 3), full first-order logic (Sect. 4), and first-order logic with linear integer arithmetic (Sect. 5). Semantics of any such logic \mathcal{L} is defined by identifying a class $\mathcal{S}_{\mathcal{L}}$ of structures (U, I) , where U is a non-empty *universe*, and I is an *interpretation* that maps predicates $p \in P$ to relations over U , functions $f \in F$ to set-theoretic functions over U , and constants $c \in C$ to values in U . Given (U, I) , the evaluation of terms and formulae is defined recursively as is common. A closed formula is called *valid* if it evaluates to *true* for all structures $(U, I) \in \mathcal{S}_{\mathcal{L}}$, and *satisfiable* if it evaluates to *true* for at least one structure.

2.2 Sequent Calculi with Constraints

Throughout the paper we will work with the *constraint sequent calculus* that is introduced in [23]. The calculus differs from normal Gentzen-style sequent calculi [11] in that every sequent $\Gamma \vdash \Delta$ is annotated with a constraint C (written $\Gamma \vdash \Delta \Downarrow C$) that captures unification conditions derived in a sub-proof. Such unification conditions come into play when free variables (which technically are treated as constants) are used to instantiate quantified formulae. All calculi in this paper are designed such that constraints cannot contain uninterpreted predicates or functions, so that validity/satisfiability of constraints is decidable. Proof procedures and refinements for the calculi are discussed in [23, 22].

More formally, if Γ, Δ are finite sets of closed formulae (the *antecedent* and *succedent*) and C is a closed formula, then $\Gamma \vdash \Delta \Downarrow C$ is called a *constrained sequent*. A sequent $\Gamma \vdash \Delta \Downarrow C$ is called *valid* if the formula $(\bigwedge \Gamma \wedge C) \rightarrow \bigvee \Delta$ is valid. A calculus rule is a binary relation between finite sets of sequents (the premises) and single sequents (the conclusion). Proof trees are defined as is common as trees growing upwards in which each node is labelled with a constrained sequent, and in which each node that is not a leaf is related with the nodes directly above through an instance of a calculus rule. A proof is closed if it is finite, and if all leaves are justified by a rule instance without premises.

2.3 The Basic Calculus for Function-Free First-Order Logic

At the core of all calculi introduced in this paper is a calculus for first-order logic with equality, at this point including uninterpreted predicates, but no functions:

$$\phi_{\text{FOL}} ::= \phi \wedge \phi \mid \phi \vee \phi \mid \neg\phi \mid \forall x.\phi \mid \exists x.\phi \mid s \doteq s \mid p(\bar{s}) \quad s ::= c \mid x$$

Since functions and arithmetic are not included in the logic, terms can only be (symbolic) constants or bound variables. Semantics is defined over the class \mathcal{S}_{FOL} of structures (U, I) with arbitrary non-empty universe U . The constraint calculus PredEq^C for the logic is shown in Fig. 1, with constraints consisting of (possibly negated) equalities, Boolean connectives, and quantifiers. The validity

$$\begin{array}{c}
\frac{\Gamma, \phi \vdash \Delta \Downarrow C \quad \Gamma, \psi \vdash \Delta \Downarrow D}{\Gamma, \phi \vee \psi \vdash \Delta \Downarrow C \wedge D} \vee L \quad \frac{\Gamma, \phi, \psi \vdash \Delta \Downarrow C}{\Gamma, \phi \wedge \psi \vdash \Delta \Downarrow C} \wedge L \quad \frac{\Gamma \vdash \phi, \Delta \Downarrow C}{\Gamma, \neg \phi \vdash \Delta \Downarrow C} \neg L \\
\frac{\Gamma \vdash \phi, \Delta \Downarrow C \quad \Gamma \vdash \psi, \Delta \Downarrow D}{\Gamma \vdash \phi \wedge \psi, \Delta \Downarrow C \wedge D} \wedge R \quad \frac{\Gamma \vdash \phi, \psi, \Delta \Downarrow C}{\Gamma \vdash \phi \vee \psi, \Delta \Downarrow C} \vee R \quad \frac{\Gamma, \phi \vdash \Delta \Downarrow C}{\Gamma \vdash \neg \phi, \Delta \Downarrow C} \neg R \\
\frac{\Gamma, [x/c]\phi, \forall x. \phi \vdash \Delta \Downarrow [x/c]C}{\Gamma, \forall x. \phi \vdash \Delta \Downarrow \exists x. C} \forall L \quad \frac{\Gamma, [x/c]\phi \vdash \Delta \Downarrow [x/c]C}{\Gamma, \exists x. \phi \vdash \Delta \Downarrow \forall x. C} \exists L \quad \frac{\Gamma \vdash \Delta \Downarrow C}{\Gamma, s \doteq t \vdash \Delta \Downarrow s \not\equiv t \vee C} =L \\
\frac{\Gamma \vdash [x/c]\phi, \exists x. \phi, \Delta \Downarrow [x/c]C}{\Gamma \vdash \exists x. \phi, \Delta \Downarrow \exists x. C} \exists R \quad \frac{\Gamma \vdash [x/c]\phi, \Delta \Downarrow [x/c]C}{\Gamma \vdash \forall x. \phi, \Delta \Downarrow \forall x. C} \forall R \quad \frac{*}{\Gamma \vdash s \doteq t, \Delta \Downarrow s \doteq t} =R \\
\frac{*}{\Gamma, p(s_1, \dots, s_n) \vdash p(t_1, \dots, t_n), \Delta \Downarrow \bigwedge_i s_i \doteq t_i} \text{PC} \quad \frac{[s/t]\Gamma, s \doteq t \vdash [s/t]\Delta \Downarrow C}{\Gamma, s \doteq t \vdash \Delta \Downarrow C} =\text{RED}
\end{array}$$

Fig. 1. The rules of the calculus PredEq^C for first-order predicate logic. In all rules, c is a constant that does not occur in the conclusion: in contrast to the use of Skolem functions and free variables in tableaux, the same kinds of symbols (constants) are used to handle both existential and universal quantifiers. Arbitrary renaming of bound variables is allowed in the constraints when necessary to avoid variable capture.

of formulae of this kind is decidable by quantifier elimination [11]. The calculus is analytic and contains two rules for each formula constructor, as well as a closure rule PC to unify complementary literals. As an optimisation, the rule =RED can be used to destructively apply equations; the rule is not necessary to establish completeness, but relevant (together with further refinements) to turn PredEq^C into a practical calculus [23, 22].

Lemma 1 (Soundness [22]). *If a sequent $\Gamma \vdash \Delta \Downarrow C$ is provable in PredEq^C , then it is valid (holds in all \mathcal{S}_{FOL} -structures).*

In particular, proving a sequent $\Gamma \vdash \Delta \Downarrow C$ with a valid constraint C implies that also the implication $\bigwedge \Gamma \rightarrow \bigvee \Delta$ is valid. This gives rise to a constraint-based proof procedure that iteratively constructs proof trees for an input sequent $\Gamma \vdash \Delta \Downarrow ?$ with a yet unknown constraint. The constraints in a proof can be filled in once all proof branches have been closed. In each iteration, the procedure checks whether the constraint generated by the current proof is valid, in which case the procedure can terminate with the result that the input problem has been proven; otherwise, the current proof has to be unfolded further. Strategies for generating proofs (without the need for backtracking, i.e., undoing previous proof steps) are discussed in [23].

Example 2. We show how to prove $\neg \forall x. (\neg p(x) \vee x \doteq c) \vee \neg p(d) \vee p(c)$, in which $p \in P$ is a unary predicate and $c, d \in C$ are constants:

$$\begin{array}{c}
\frac{*}{\overline{p(d) \vdash p(a) \Downarrow d \doteq a}} \text{PC} \quad \frac{*}{\overline{a \doteq c, p(d) \vdash p(c) \Downarrow d \doteq c}} \text{PC} \\
\frac{\overline{p(d) \vdash p(a) \Downarrow d \doteq a} \quad \overline{a \doteq c, p(d) \vdash p(c) \Downarrow d \doteq c}}{\overline{\neg p(a), p(d) \vdash \dots \Downarrow d \doteq a} \quad \overline{a \doteq c, p(d) \vdash p(c) \Downarrow a \not\equiv c \vee d \doteq c}} =L \\
\frac{\dots, \overline{\neg p(a) \vee a \doteq c, p(d) \vdash p(c) \Downarrow d \doteq a \wedge (a \not\equiv c \vee d \doteq c)}}{\overline{\forall x. (\neg p(x) \vee x \doteq c), p(d) \vdash p(c) \Downarrow R}} \vee L \\
\frac{\overline{\forall x. (\neg p(x) \vee x \doteq c), p(d) \vdash p(c) \Downarrow R}}{\vdash \overline{\neg \forall x. (\neg p(x) \vee x \doteq c) \vee \neg p(d) \vee p(c)} \Downarrow R} \vee R^*, \neg R^*
\end{array}$$

In order to instantiate the universal quantifier, the fresh constant a is introduced; the constant is quantified existentially in the derived constraints, and therefore can be seen as a “free variable.” The constraints on the right-hand side of \Downarrow are practically filled in *after* closing the proof using PC. The validity of the original formula follows from the validity of $R = \exists x.(d \doteq x \wedge (x \neq c \vee d \doteq c))$.

Lemma 3 (Completeness [24]). *Suppose ϕ is closed and valid. Then there is a valid constraint C such that $\vdash \phi \Downarrow C$ is provable in PredEq^C .*

3 Positive Unit Hyper-Resolution

As argued in Sect. 1.1, axioms and quantified formulae (in particular in verification problems) are often manually formulated with a clear, directed application strategy in mind. This makes it possible to systematically instantiate axioms in a manner that more resembles the execution of a functional or logic program than the search for a proof. From a practical point of view, providing support for this style of reasoning (even if it is only applicable to a subset of input problems) is crucial to achieve the scalability needed for applications. We integrate such user-guided reasoning into our calculus with the help of concepts from the *positive unit hyper-resolution* (PUHR) calculus, an approach first used in the SATCHMO theorem prover [14, 15]. PUHR will be used in Sect. 4 to simulate the e-matching method common in SMT solvers.

PUHR is a tableau procedure in which clauses are instantiated by matching negative literals on (ground) literals already present on a proof branch. Starting from the calculus PredEq^C defined in the last section, we introduce a similar rule in our hyper-resolution sequent calculus PredEqHR^C , instantiating quantified formulae that are “guarded” by negative literals $\neg p_1(\bar{t}_1), \dots, \neg p_n(\bar{t}_n)$ using symbols from matching literals $p_1(\bar{s}_1), \dots, p_n(\bar{s}_n)$ in the antecedent of a sequent:

$$\frac{\Gamma, \{p_i(\bar{s}_i)\}_{i=1}^n, \forall \bar{x}. (\bigvee_{i=1}^n \neg p_i(\bar{t}_i) \vee \phi), \text{simp}(\forall \bar{x}. (\bigvee_{i=1}^n \bar{s}_i \neq \bar{t}_i \vee \phi)) \vdash \Delta \Downarrow C}{\Gamma, \{p_i(\bar{s}_i)\}_{i=1}^n, \forall \bar{x}. (\bigvee_{i=1}^n \neg p_i(\bar{t}_i) \vee \phi) \vdash \Delta \Downarrow C} \forall\text{L-M}$$

Given literals $\{p_i(\bar{s}_i)\}_{i=1}^n$ in a sequent, a quantified formula $\forall \bar{x}. (\bigvee_{i=1}^n \neg p_i(\bar{t}_i) \vee \phi)$ can be instantiated using the argument terms \bar{s}_i by simultaneously solving the systems $\bar{s}_i \doteq \bar{t}_i$ of equalities. In contrast to the original PUHR [14], we do not require formulae to be range restricted. Note that the formula ϕ might be *false* and disappear, and that the literals $\{p_i(\bar{s}_i)\}_{i=1}^n$ are not necessarily distinct. The solving of equalities is formulated using a recursive simplification function *simp*:

$$\begin{aligned} \text{simp}(\forall \bar{x}. (t \neq t \vee \phi)) &= \text{simp}(\forall \bar{x}. \phi) \\ \text{simp}(\forall \bar{x}. (x_i \neq t \vee \phi)) &= \text{simp}(\forall \bar{x}. [x_i/t]\phi) && (x_i \neq t) \\ \text{simp}(\forall \bar{x}. (t \neq x_i \vee \phi)) &= \text{simp}(\forall \bar{x}. [x_i/t]\phi) && (x_i \neq t) \\ \text{simp}(\forall \bar{x}. (s \neq t \vee \phi)) &= s \neq t \vee \text{simp}(\forall \bar{x}. \phi) && (s, t \notin \bar{x}) \\ \text{simp}(\forall \bar{x}. \phi) &= \forall (\bar{x} \cap \text{fv}(\phi)). \phi && (\text{otherwise}) \end{aligned}$$

A rule $\exists\text{R-M}$ similar to $\forall\text{L-M}$ is introduced for existentially quantified formulae $\exists \bar{x}. (\bigwedge_{i=1}^n p_i(\bar{t}_i) \wedge \phi)$ in the succedent. The soundness of the new rules is

4.1 Relational Encoding of Functions

We consider first-order logic including function symbols, which means that the grammar for terms shown in the beginning of Sect. 2.3 is extended to:

$$s ::= c \mid x \mid f(s, \dots, s)$$

where $f \in F$ ranges over function symbols. For the purpose of the encoding of functions into relations, we assume that a fresh $(n + 1)$ -ary uninterpreted predicate $f_p \in P$ exists for every n -ary uninterpreted function $f \in F$, representing the graph of f . The relation f_p satisfies two axioms, *functionality* and *totality*:

$$Fun_f = \forall \bar{x}, y_1, y_2. (\neg f_p(\bar{x}, y_1) \vee \neg f_p(\bar{x}, y_2) \vee y_1 = y_2), \quad Tot_f = \forall \bar{x}. \exists y. f_p(\bar{x}, y).$$

We can then translate from formulae ϕ over the functional vocabulary F (and relational vocabulary P) to formulae ϕ_{Rel} purely over the relational vocabulary P . This can be done by means of the following rewriting rules:

$$\begin{aligned} \exists\text{-enc:} \quad & \psi[f(\bar{t})] \rightsquigarrow \exists x. (f_p(\bar{t}, x) \wedge \psi[x]) \\ \forall\text{-enc:} \quad & \psi[f(\bar{t})] \rightsquigarrow \forall x. (\neg f_p(\bar{t}, x) \vee \psi[x]) \end{aligned}$$

Both rules have the side condition that rewritten occurrences of $f(\bar{t})$ must not be in the scope of quantifiers binding variables in the terms \bar{t} ; furthermore, the variable x must be fresh in $\psi[f(\bar{t})]$. It is possible, however, to apply the rewriting rules to arbitrary sub-formulae of a given formula ϕ ; in other words, the predicate and quantifier that encode a function application $f(\bar{t})$ can be placed arbitrarily in the rewritten formula, as long as the function application remains in the scope of the quantifier. Rewriting strategies are discussed later in this section.

Lemma 6. *Suppose ϕ is a closed formula over the vocabulary F , and ϕ_{Rel} is a function-free formula obtained from ϕ by application of the rewriting rules \exists -enc and \forall -enc. Then ϕ is valid iff $\bigwedge_{f \in F} (Fun_f \wedge Tot_f) \rightarrow \phi_{Rel}$ is valid.*

Since the calculi $PredEq^C$ and $PredEqHR^C$ are sound and complete for first-order logic without function symbols, we can therefore construct calculi for first-order logic including functions by first encoding functions as relations.

4.2 Ground Reasoning and Congruence Closure

We first concentrate on quantifier-free first-order formulae with functions. In this setting, it is easy to see that the hyper-resolution calculus $PredEqHR^C$, in combination with the functionality axioms Fun_f for functions f , is able to simulate congruence closure procedures. This is supported by the following strengthened version of Lem. 6, which observes that totality axioms are not necessary when solving essentially ground formulae:

Lemma 7. *Suppose ϕ is a closed formula over the vocabulary F , and ϕ_{Rel} a function-free formula obtained from ϕ by application of the rewriting rules \exists -enc and \forall -enc that contains \forall -quantifiers only in positive positions, and \exists -quantifiers only in negative positions. Then ϕ is valid iff $\bigwedge_{f \in F} Fun_f \rightarrow \phi_{Rel}$ is valid.*

Example 9. Consider the quantified formula $\forall x.f(x) \doteq g(x)$. Four possible ways of encoding the formula using relations, corresponding to different strategies when applying the rules $\forall\text{-enc}$ and $\exists\text{-enc}$, are:

$$\forall x.\exists y, z.(f_p(x, y) \wedge g_p(x, z) \wedge y \doteq z) \quad (7)$$

$$\forall x, y.(\neg f_p(x, y) \vee \exists z.(g_p(x, z) \wedge y \doteq z)) \quad (8)$$

$$\forall x, z.(\neg g_p(x, z) \vee \exists y.(f_p(x, y) \wedge y \doteq z)) \quad (9)$$

$$\forall x, y, z.(\neg f_p(x, y) \vee \neg g_p(x, z) \vee y \doteq z) \quad (10)$$

Each of the relational formulae corresponds to a particular selection of triggers in $\forall x.f(x) \doteq g(x)$:

- in (7), no triggers have been chosen, with the result that the hyper-resolution rule $\forall\text{L-M}$ is not applicable. Instantiation of (7) is only possible using the rule $\forall\text{L}$, replacing the bound variable x with an existentially quantified constant that can later unified with some term.
- in (8), the term $f(x)$ (corresponding to the negative literal $f_p(x, y)$) has been selected as trigger. In the calculus PredEqHR^C , (8) can only be instantiated using the rule $\forall\text{L-M}$, and only in case a literal $f_p(s, t)$ occurs in the antecedent of a sequent, substituting the terms s, t for the variables x, y . This corresponds to e-matching the expression $f(x)$ on a node $f(t)$ of a congruence graph. No free variables are needed to instantiate (8).
- similarly, in (9) the term $g(x)$ is trigger.
- in (10), both $f(x)$ and $g(x)$ have been chosen as a *multi-trigger*, which means that (10) only can be instantiated if literals $f_p(s, t)$ and $g_p(s', t')$ occur in an antecedent. In this case, the instance $s \neq s' \vee t \doteq t'$ will be generated, expressing that the equality $t \doteq t'$ can be assumed if s and s' are unifiable. In terms of e-graphs, the formula would only be instantiated if the e-graph contains nodes $f(s), g(s')$ such that s, s' are in the same equivalence class.

The following proof fragment illustrates how (9) can be instantiated referring to a literal $g_p(a, b)$ in the antecedent, effectively adding $f_p(a, b)$ to the sequent:

$$\frac{\frac{\frac{g_p(a, b), (9), f_p(a, b), b \doteq u \vdash \Downarrow [y/u]C}{g_p(a, b), (9), f_p(a, u), u \doteq b \vdash \Downarrow [y/u]C} =_{\text{RED}}}{g_p(a, b), (9), \exists y.(f_p(a, y) \wedge y \doteq b) \vdash \Downarrow \forall y.C} \exists\text{L}, \wedge\text{L}}{g_p(a, b), (9) \vdash \Downarrow \forall y.C} \forall\text{L-M} \quad (11)$$

The way in which a formula ϕ is translated to ϕ_{Rel} determines how quantified sub-formulae are instantiated, in the same way as SMT solvers can be guided by specifying triggers (Alg. 1 shows how the translation can be done systematically, for a given set of triggers). However, it can be observed that the four encodings (7)–(10) are all equivalent w.r.t. provability of theorems: in combination with the axioms $\text{Fun}_f, \text{Fun}_g, \text{Tot}_f, \text{Tot}_g$ each of the formulae can simulate each other formula. *The choice of triggers in formulae therefore only influences efficiency, not completeness.* For instance, formula (9) in (11) can be replaced

Algorithm 1: ENCODETRIGGER: relational encoding of a quantified formula for a specific set of triggers

Input: Formula $\forall \bar{x}.\phi$, set T of trigger terms with variables from \bar{x}
Output: Relational formula ϕ_{Rel}

$qvars \leftarrow \{x \mid x \in \bar{x}\};$
 $premises \leftarrow \emptyset;$
while T contains function terms **do**
 pick (sub)term $f(\bar{t})$ in T s.t. \bar{t} does not contain functions;
 pick fresh variable y ;
 $qvars \leftarrow qvars \cup \{y\};$
 $premises \leftarrow premises \cup \{f_p(\bar{t}, y)\};$
 substitute y for $f(\bar{t})$ everywhere in T and ϕ ;
end
apply \exists -enc exhaustively to ϕ ;
return $\forall_{x \in qvars} . (\bigvee_{p \in premises} \neg p \vee \phi);$

with (8) in the following way (the constraints of the sequents have been left out for sake of brevity):

$$\begin{array}{c}
\frac{*}{\dots \vdash x \doteq a} =R \quad \frac{f_p(x, b), g_p(x, b), g_p(a, b), v \doteq b \vdash}{f_p(x, v), g_p(x, v), g_p(a, b), v \doteq b \vdash} =RED \\
\frac{\dots, f_p(x, v), g_p(x, v), g_p(a, b), x \not\equiv a \vee v \doteq b \vdash}{\text{Fun}_g, f_p(x, v), g_p(x, v), v \doteq u, g_p(a, b) \vdash} \forall L, \neg L \\
\frac{\text{Fun}_g, f_p(x, v), g_p(x, v), v \doteq u, g_p(a, b) \vdash}{\text{Fun}_g, f_p(x, u), g_p(x, v), u \doteq v, g_p(a, b) \vdash} =RED \\
\frac{\text{Fun}_g, f_p(x, u), g_p(x, v), u \doteq v, g_p(a, b) \vdash}{\text{Fun}_g, \dots, f_p(x, u), \exists z. (g_p(x, z) \wedge u \doteq z), g_p(a, b) \vdash} \exists L, \wedge L \\
\frac{\text{Fun}_g, \dots, f_p(x, u), \exists z. (g_p(x, z) \wedge u \doteq z), g_p(a, b) \vdash}{\text{Fun}_g, Tot_f, f_p(x, u), g_p(a, b), (8) \vdash} \forall L, \neg L \\
\frac{\text{Fun}_g, Tot_f, f_p(x, u), g_p(a, b), (8) \vdash}{\text{Fun}_g, Tot_f, g_p(a, b), (8) \vdash} \forall L, \exists L
\end{array}$$

This illustrates that PUHR/e-matching-based reasoning (through $\forall L$ -M and $\exists R$ -M) can be mixed freely with free variable reasoning (through $\forall L$ and $\exists R$). Proofs constructed without applying the rules $\forall L$ and $\exists R$ closely correspond to the ground reasoning in an SMT solver, while each application of $\forall L$ or $\exists R$ conceptually introduces a free variable that, at a later point during proof construction, can be unified with other terms, extracting unification conditions in the form of constraints.

5 Extension to Linear Integer Arithmetic

All techniques discussed so far carry over to first-order logic modulo the theory of linear integer arithmetic (FOL(LIA)), via integration into the calculus defined in [23]. The syntax of FOL(LIA) is defined by the grammar in the beginning of Sect. 2.1 and combines first-order logic (with uninterpreted predicates and functions) with arithmetic terms and predicates. Semantics is defined over structures (\mathbb{Z}, I) with the set of integers as universe.

$$\begin{array}{c}
\frac{\Gamma, t \doteq 0 \vdash \phi[s + \alpha \cdot t], \Delta \Downarrow C}{\Gamma, t \doteq 0 \vdash \phi[s], \Delta \Downarrow C} =_{\text{RED-Z}} \quad \frac{\Gamma, s \leq 0, t \leq 0, \alpha s + \beta t \leq 0 \vdash \Delta \Downarrow C}{\Gamma, s \leq 0, t \leq 0 \vdash \Delta \Downarrow C} \leq_{\text{L-Z}} \\
\frac{\Gamma, p(s_1, \dots, s_n) \vdash p(t_1, \dots, t_n), \bigwedge_i s_i - t_i \doteq 0, \Delta \Downarrow C}{\Gamma, p(s_1, \dots, s_n) \vdash p(t_1, \dots, t_n), \Delta \Downarrow C} \text{PU-Z} \\
\frac{*}{\Gamma, \phi_1, \dots, \phi_n \vdash \psi_1, \dots, \psi_m, \Delta \Downarrow \neg\phi_1 \vee \dots \vee \psi_1 \vee \dots} \text{CLOSE}
\end{array}$$

Fig. 2. A selection of rules of the calculus PresPred^C ; for a complete list see [23]. In $=_{\text{RED-Z}}$, α is a literal; we write $\phi[s]$ in the succedent to denote that s occurs in an arbitrary formula in the sequent, which can in particular also be in the antecedent. In $\leq_{\text{L-Z}}$, $\alpha, \beta > 0$ are positive literals. In CLOSE-Z , the formulae $\phi_1, \dots, \phi_n, \psi_1, \dots, \psi_m$ do not contain uninterpreted predicates.

As for FOL, we first introduce a calculus for the function-free fragment of FOL(LIA). The integration of functions is then done in the same way as in Sect. 3, 4 with the help of a relational encoding. The calculus PresPred^C for the function-free fragment consists of the rules in Fig. 1, together with a number of rules specific for linear integer arithmetic, a selection of which are shown in Fig. 2 (as a result, the rules $=_{\text{L}}$, $=_{\text{R}}$, $=_{\text{RED}}$, and PC of the first-order calculus can be removed); in the full calculus, also simplification and splitting rules are needed [23]. A more general closure rule CLOSE has to be used than in PredEq^C to support disjunctive constraints. Constraints in PresPred^C are always formulae in Presburger arithmetic (PA), i.e., do not contain uninterpreted predicates.

Lemma 10 (Soundness [23]). *If a sequent $\Gamma \vdash \Delta \Downarrow C$ can be proven in PresPred^C , then it is valid.*

The logic FOL(LIA) subsumes Presburger arithmetic. Since the logic of quantified Presburger arithmetic with predicates is Π_1^1 -complete [10], no complete calculi can exist for FOL(LIA); however, it can be shown that the calculi introduced in this section are complete for relevant and non-trivial fragments:

Lemma 11 (Completeness [23]). *Suppose ϕ is a closed formula without functions or constants in one of the following fragments:*

- (i) ϕ does not contain uninterpreted predicates (i.e., in Presburger arithmetic);
- (ii) ϕ contains universal (exist.) quantifiers only in positive (negative) positions;
- (iii) ϕ contains universal (exist.) quantifiers only in negative (positive) positions;
- (iv) ϕ is of the form $\forall \bar{x}.(\sigma \rightarrow \psi)$, where σ is a formula in Presburger arithmetic (without uninterpreted predicates) that has only finitely many solutions in \bar{x} , and ψ contains universal (existential) quantifiers only in negative (positive) positions (i.e., a formula accepted by the $\mathcal{ME}(LIA)$ calculus [4]).

Then there is a valid constraint C such that $\vdash \phi \Downarrow C$ is provable in PresPred^C .

Practically, it can be observed that PresPred^C can often also be applied successfully to formulae outside of those fragments.

5.1 Hyper-Resolution and E-Matching for FOL(LIA)

The unit hyper-resolution rule $\forall\text{L-M}$ (and similarly the rule $\exists\text{R-M}$) defined in Sect. 3 can be integrated in the calculus PresPred^C in the same way as in the earlier first-order calculi, in order to instantiate formulae $\forall\bar{x}. (\bigvee_{i=1}^n \neg p_i(\bar{t}_i) \vee \phi)$ by matching. In this context, the simplification function *simp* can be (but does not have to be) replaced with a function tailored to integer arithmetic, i.e., a function that is able to solve the system $\bigvee_{i=1}^n \bar{s}_i \neq \bar{t}_i$ modulo integer arithmetic.

The calculus PresPredHR^C is derived from PresPred^C by adding the rules $\forall\text{L-M}$ and $\exists\text{R-M}$, and by imposing the side condition that the rule $\forall\text{L}$ is no longer applied to formulae of the shape $\forall\bar{x}. (\bigvee_{i=1}^n \neg p_i(\bar{t}_i) \vee \phi)$; similarly for the rule $\exists\text{R}$. As before, the soundness of the rules $\forall\text{L-M}$ and $\exists\text{R-M}$ is immediate. We can also observe that PresPredHR^C is relatively complete, in the sense that formulae that are provable in PresPred^C can also be proven using PresPredHR^C :

Lemma 12. *Suppose $\Gamma \vdash \Delta \Downarrow C$ is provable in PresPred^C , where C is valid. Then there is a valid constraint C' so that PresPredHR^C can prove $\Gamma \vdash \Delta \Downarrow C'$.*

Encoding of functions. The relational encoding of functions from Sect. 4 can be used to obtain a calculus for the full logic FOL(LIA) with functions. Although there are no complete calculi for the full logic, we can observe that PresPred^C (and therefore, by Lem. 12, PresPredHR^C) can handle at least all formulae that can be proven by considering a finite set of ground instances:

Lemma 13. *Suppose $\exists\bar{x}.\phi$ is a closed formula in FOL(LIA), with functions taken from a finite set F , such that ϕ is quantifier-free. If there is a valid disjunction $\bigvee_{i=1}^n [\bar{x}/\bar{t}_i]\phi$ of ground instances of $\exists\bar{x}.\phi$, then there is a valid constraint C such that $\{\text{Fun}_f, \text{Tot}_f\}_{f \in F} \vdash (\exists\bar{x}.\phi)_{\text{Rel}} \Downarrow C$ is provable in PresPred^C .*

The lemma directly generalises to disjunctions of existentially quantified formulae, which in particular entails that PresPred^C is complete for the class of *essentially uninterpreted formulae* F (modulo linear integer arithmetic) with finite ground instantiation F^* defined in [9], and thus also for the array property fragment [6] (PresPred^C cannot easily be turned into a decision procedure, however, since it would be unclear how to ensure termination on invalid problems).

6 Experiments and Related Work

We have implemented the described calculus PresPredHR^C for FOL(LIA) in the theorem prover PRINCESS,² and are in the process of adding further optimisations. PRINCESS uses the relational encoding from Sect. 4 to represent functions, and heuristics similar to the ones in Simplify [7] to automatically identify triggers in quantified formulae; redundancy criteria [22] and theory propagation help to reduce the number of instances generated from quantified formulae. PRINCESS is able to handle all of the examples discussed in Sect. 1.1.

² <http://www.philipp.ruemmer.org/princess.shtml>

	AUFLIA+p (193)	AUFLIA-p (193)
Z3	191	191
Princess	145	137
CVC3	132	128

Fig. 3. Number of solved benchmarks, out of 2×193 unsatisfiable (scrambled) AUFLIA benchmarks selected in the SMT competition 2011. Experiments with PRINCESS were done on an Intel Core i5 2-core machine with 3.2GHz, with a timeout of 1200s, heap-space limited to 4Gb. The benchmarks in AUFLIA+p contain hand-written triggers for most of the quantified formulae, while all triggers have been removed in AUFLIA-p. The corresponding figures for Z3 and CVC3 are the results obtained during the SMT competition 2011 (<http://www.smtexec.org/exec/?jobs=856>).

To evaluate the overhead of the relational function encoding, we compared the performance of PRINCESS with the SMT solvers CVC3 [3] and Z3 [19], using benchmarks selected in the SMT competition 2011. Since our work concentrates on proof construction, we only considered unsatisfiable benchmarks, removing 13 satisfiable AUFLIA problems in each category. The results show that PRINCESS, while currently not being able to compete with the fastest SMT solver Z3, performs better than the (state-of-the-art) e-matching-based CVC3. This is promising, since PRINCESS does not (yet) use SMT techniques like lemma learning, which are important for large or propositionally complex problems. PRINCESS can solve most benchmarks using e-matching alone, but uses free variables in 17 of the (solved) benchmarks, typically in smaller (but harder) instances.

Related Work E-matching is today used in most SMT solvers, based on techniques that go back to the Simplify prover [7] and The Stanford Pascal Verifier [20]; since then, various refinements of the e-matching approach have been published, for instance [8, 18]. To the best of our knowledge, e-matching has not previously been combined with free variable methods. An instantiation method similar to e-matching, but with much stronger completeness results, has been published in [9] and is used in Z3; a comparison with our method is in Sect. 5.1.

There is a large body of work on integrating theories into resolution and superposition calculi (e.g., [25, 2, 13, 1]), as well as on the integration of resolution into SMT [17]. These approaches completely avoid e-matching, offering stronger completeness guarantees but limiting the possibility of user-provided guidance.

The model evolution calculus has been extended to theories, including integer arithmetic [4, 5]. Our approach resembles model evolution in that it also uses free variables in a tableaux setting, albeit in a more “rigid”/global manner. Further differences are that $\mathcal{ME}(LIA)$ works on clauses, only supports a restricted form of existential quantification, and has a more explicit representation of models.

References

1. Althaus, E., Kruglov, E., Weidenbach, C.: Superposition modulo linear arithmetic sup(la). In: FroCos. LNCS, vol. 5749, pp. 84–99. Springer (2009)

2. Bachmair, L., Ganzinger, H., Waldmann, U.: Refutational theorem proving for hierarchic first-order theories. *Appl. Algebra Eng. Commun. Comput.* 5 (1994)
3. Barrett, C., Tinelli, C.: CVC3. In: *Proceedings, CAV*. LNCS, vol. 4590, pp. 298–302. Springer (Jul 2007)
4. Baumgartner, P., Fuchs, A., Tinelli, C.: ME(LIA) – Model Evolution With Linear Integer Arithmetic Constraints. In: *LPAR*. LNCS, vol. 5330. Springer (2008)
5. Baumgartner, P., Tinelli, C.: Model evolution with equality modulo built-in theories. In: *CADE*. LNCS, vol. 6803, pp. 85–100. Springer (2011)
6. Bradley, A.R., Manna, Z., Sipma, H.B.: What’s decidable about arrays? In: *VMCAI*. LNCS, vol. 3855, pp. 427–442. Springer (2006)
7. Detlefs, D., Nelson, G., Saxe, J.B.: Simplify: A theorem prover for program checking. *Journal of the ACM* 52(3) (2005)
8. Ge, Y., Barrett, C., Tinelli, C.: Solving quantified verification conditions using satisfiability modulo theories. In: *CADE*. LNCS, vol. 4603. Springer (2007)
9. Ge, Y., de Moura, L.M.: Complete instantiation for quantified formulas in satisfiability modulo theories. In: *CAV*. pp. 306–320 (2009)
10. Halpern, J.Y.: Presburger arithmetic with unary predicates is Π_1^1 complete. *Journal of Symbolic Logic* 56 (1991)
11. Harrison, J.: *Handbook of Practical Logic and Automated Reasoning*. Cambridge University Press (2009)
12. Klebanov, V., Müller, P., Shankar, N., Leavens, G.T., Wüstholtz, V., Alkassar, E., Arthan, R., Bronish, D., Chapman, R., Cohen, E., Hillebrand, M., Jacobs, B., Leino, K.R.M., Monahan, R., Piessens, F., Polikarpova, N., Ridge, T., Smans, J., Tobies, S., Tuerk, T., Ulbrich, M., Weiß, B.: *The 1st Verified Software Competition: Extended experience report* (2011)
13. Korovin, K., Voronkov, A.: Integrating linear arithmetic into superposition calculus. In: *CSL*. LNCS, vol. 4646, pp. 223–237. Springer (2007)
14. Manthey, R., Bry, F.: A hyperresolution-based proof procedure and its implementation in Prolog. In: *GWAI*. pp. 221–230. Springer (1987)
15. Manthey, R., Bry, F.: SATCHMO: A theorem prover implemented in Prolog. In: *Proceedings, CADE*. pp. 415–434. LNCS, Springer (1988)
16. McCarthy, J.: Towards a mathematical science of computation. In: Popplewell, C.M. (ed.) *Information Processing 1962*. pp. 21–28. North Holland (1963)
17. de Moura, L., Bjørner, N.: Engineering DPLL(T) + saturation. In: *Proceedings, IJCAR*. LNCS, vol. 5195. Springer (2008)
18. de Moura, L.M., Bjørner, N.: Efficient e-matching for SMT solvers. In: *Proceedings, CADE*. pp. 183–198. LNCS, Springer (2007)
19. de Moura, L.M., Bjørner, N.: Z3: An efficient SMT solver. In: *TACAS*. LNCS, vol. 4963, pp. 337–340. Springer (2008)
20. Nelson, G.: *Techniques for program verification*. Tech. Rep. CSL-81-10, Xerox Palo Alto Research Center (1981)
21. Nelson, G., Oppen, D.C.: Fast decision procedures based on congruence closure. *J. ACM* 27, 356–364 (April 1980)
22. Rümmer, P.: *Calculi for Program Incorrectness and Arithmetic*. Ph.D. thesis, University of Gothenburg (2008)
23. Rümmer, P.: A constraint sequent calculus for first-order logic with linear integer arithmetic. In: *LPAR*. LNCS, Springer (2008)
24. Rümmer, P.: E-matching with free variables. Tech. rep. (2012), to appear
25. Stickel, M.E.: Automated deduction by theory resolution. *Journal of Automated Reasoning* 1(4), 333–355 (1985)