

Equality Preprocessing in Connection Calculi

Benjamin E. Oliver, Jens Otten

University of Oslo, Oslo, Norway

Abstract

Equality is a fundamental concept in first-order reasoning, yet for connection based proof methods a notoriously challenging one to handle efficiently. While paramodulation is a popular technique for dealing with equality in resolution and related calculi, there is no single practical successful solution for connection based approaches. We present an extensible system for equality preprocessing in connection calculi (EPICC) that can be used as a tool in reducing the search space of problems that contain equality. We specify a number of preprocessing rules, describe an implementation of these rules and compare it with existing approaches for dealing with equality in connection calculi.

Keywords

Automated Reasoning, First-order Logic, Equality, Connection Calculus

1. Introduction

Equality is an essential concept when reasoning about everyday life. Knowing that London is the capital of the *United Kingdom* and that the *United Kingdom* and the *UK* are equal, i.e. they denote the same country, lets us deduce that London is the capital of the *UK*. Indeed, it is not surprising that the largest collection of problems for automated reasoning or, more specifically, automated theorem proving (ATP), the TPTP library [1], uses equality in the formalisation of a vast number of problems. Yet, for a concept so ubiquitous its automation is not straightforward.

In mathematics and first-order logic, equality is usually represented by “=” and the above statement could (in short) be encoded using the first-order formula

$$(P(a) \wedge a = b) \Rightarrow P(b)$$

In order to solve this problem, ATP systems have to incorporate techniques for equality. While paramodulation [2] is a successful technique for dealing with equality in the popular resolution method [3], the situation is more complicated for tableau or connection calculi [4, 5]. For example, rigid E-unification [6] is not decidable and its use practically infeasible due to its complexity [7]. A more restricted technique called bounded rigid E-unification has been implemented in the tableau prover ePrincess [8, 9], but cannot easily be extended to connection calculi.

So far, the most successful technique for dealing with equality in connection calculi, as also implemented in the leanCoP prover [10, 11], adds the equality axioms and then uses

PAAR 2020: *Seventh Workshop on Practical Aspects of Automated Reasoning*, June 29–30, 2020, Paris, France

EMAIL: benjamin.e.oliver@gmail.com (B. E. Oliver); jeotten@ifi.uio.no (J. Otten)

URL: <https://github.com/beoliver/> (B. E. Oliver); <http://jens-otten.de/> (J. Otten)



© 2020 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

restricted backtracking [12] to limit the amount of redundancy caused by the equality axioms. But as observed in the yearly system competitions CASC, the relative performance of leanCoP compared to other provers on problems with equality is significantly lower than its relative performance on problems without equality [13].

In this paper we present a framework for preprocessing techniques in order to simplify problems containing equality. Even though the presented approach can be used in combination with any ATP procedure, we have implemented, tested and evaluated it in combination with the connection prover leanCoP. It is also tested against an implementation of the *modification method* [14], another well-known preprocessing technique to deal with equality in connection calculi.

In Section 2, we first present the details of the underlying matrix method. Section 3 introduces the preprocessing steps and rules. Section 4 gives details on how these preprocessing rules have been implemented and combined with leanCoP. In Section 5 this implementation is compared to the equality technique currently used in leanCoP and an implementation of the modification method. The paper concludes with a summary and brief outlook on further research in Section 6.

2. Preliminaries

In this section some basic concepts and notations are introduced, such as the matrix characterization and the standard equality axioms.

2.1. First-Order Logic and Matrix Characterization

The standard notation for first-order formulae is used. *Terms* (denoted by s, t, u, v) are built up from functions (denoted by f), constants (a, b, c) and variables (denoted by x, y, z). An *atomic formula* (denoted by A) is comprised of predicate symbols (denoted by P) and terms. A (*first-order*) *formula* (denoted by F) is built up from atomic formulae, the connectives $\neg, \wedge, \vee, \Rightarrow$, and the first-order quantifiers \forall and \exists . A *literal* L has the form A or $\neg A$. Its *complement* \bar{L} is A if L is of the form $\neg A$; otherwise \bar{L} is $\neg L$.

A formula in (disjunctive) *clausal form* has the form $\exists x_1 \dots \exists x_n (C_1 \vee \dots \vee C_n)$, where each clause C_i is a conjunction of literals L_1, \dots, L_{m_i} .¹ It is usually represented as a set of clauses $\{C_1, \dots, C_n\}$, which is called a (clausal) *matrix* \mathbf{M} . Every formula F can be translated into a validity-preserving formula F' in clausal form.

Definition 2.1 (Matrix). *A set of clauses is represented as a matrix. A matrix \mathbf{M} of a formula consists of its clauses $\{C_1, \dots, C_n\}$, in which each clause C_i is a set of its literals $\{L_1, \dots, L_{m_i}\}$.*

In the *graphical representation* of a matrix, its clauses are arranged horizontally, while the literals of each clause are arranged vertically (see Figure 1).

A *connection* $\{P(\dots), \neg P(\dots)\}$ is a set of two literals with the same predicate symbol, of which (exactly) one is negated. A *first-order* or *term substitution* σ is a mapping from the set of term

¹Even though the use of a *conjunctive clausal form* (cnf) is common, a disjunctive clausal form (dnf) is used for historical and practical reasons; the difference between both forms is marginal (a formula F in dnf is *valid* iff $\neg F$ in cnf is *unsatisfiable*).

$$\left[\begin{array}{c} [L_1^1] \\ \vdots \\ [L_i^1] \end{array} \quad \dots \quad \begin{array}{c} [L_1^n] \\ \vdots \\ [L_j^n] \end{array} \right]$$

Figure 1: Graphical representation of a matrix.

variables to the set of terms. In $\sigma(L)$ and $\sigma(C)$ all term variables x in L and C are substituted by their image $\sigma(x)$. A connection $\{L_1, L_2\}$ with $\sigma(L_1) = \sigma(L_2)$ is called σ -complementary. A path through a matrix $\mathbf{M} = \{C_1, \dots, C_n\}$ is a set of literals that contains one literal from each clause $C_i \in \mathbf{M}$, i.e. a set $\bigcup_{i=1}^n \{L'_i\}$ with $L'_i \in C_i$. The following *matrix characterization* [15] provides a simple criterion for the validity of a formula and is the basis of the connection method [5]; see also [16].

Theorem 2.1 (Matrix Characterization). *A formula F and its matrix \mathbf{M} are valid iff there exists (1) a multiplicity $\mu : \mathbf{M} \rightarrow \mathbf{N}$ (specifying the number of clause copies), (2) a term substitution σ and (3) a set of connections S , such that every path through the matrix \mathbf{M}^μ of F contains a σ -complementary connection $\{L_1, L_2\} \in S$. In \mathbf{M}^μ , clause copies have been added according to μ .*

2.2. Equality

In order to extend the language to first-order logic *with equality*, the (predefined) *equality predicate* \approx is added. Instead of $\approx (s, t)$ we use the common infix notation $s \approx t$. We also use $s \not\approx t$ as an abbreviation for $\neg(s \approx t)$. One way to specify the interpretation (or meaning) of equality is by adding equality axioms. Once these axioms have been added the equality symbols \approx and $\not\approx$ can be treated as uninterpreted predicates.

Definition 2.2 (Equality Axioms). *We will use the notation $\alpha(\mathbf{M})$ to denote the set of axiom clauses that must be generated for the matrix \mathbf{M} of a formula and $\mathbf{M} \cup \alpha(\mathbf{M})$ to indicate the resulting matrix formed by combining the original matrix and the axioms. If \mathbf{M} does not contain equality then $\alpha(\mathbf{M})$ is an empty set, however if \mathbf{M} contains equality, then $\alpha(\mathbf{M})$ is the least set such that:*

$$\left[\begin{array}{cc} [x \not\approx x] & [x \approx y] \\ & [y \not\approx x] \end{array} \quad \left[\begin{array}{c} [x \approx y] \\ [y \approx z] \\ [x \not\approx z] \end{array} \right] \right] \subseteq \alpha(\mathbf{M}) \quad (\text{reflexivity, symmetry, transitivity})$$

$$\left[\begin{array}{c} x_1 \approx y_1 \\ \vdots \\ x_n \approx y_n \\ f(x_1 \dots x_n) \not\approx f(y_1 \dots y_n) \end{array} \right] \in \alpha(\mathbf{M}) \quad (\text{for every function } f \text{ of arity } n \text{ in } \mathbf{M})$$

$$\left[\begin{array}{c} x_1 \approx y_1 \\ \vdots \\ x_n \approx y_n \\ P(x_1 \dots x_n) \\ \neg P(y_1 \dots y_n) \end{array} \right] \in \alpha(\mathbf{M}) \quad (\text{for every predicate } P \text{ of arity } n \text{ in } \mathbf{M})$$

3. Equality Preprocessing

Unlike the modification method and its derivatives, the following approach is not designed to eliminate equality entirely. Instead, it is best viewed as a set of rules that aim to balance three properties - ease of implementation, good algorithmic complexity and, finally, improvement of the performance of the proof procedure. A combination that is difficult to achieve at the same time.

3.1. Basic Notation

3.1.1. Matrix Notation

As both matrices and clauses are really nothing but sets, we will employ a special notation that allows us to focus on certain properties. The notation $\mathbf{M} = [\mathbf{C} \ S]$ is a visual representation of $\mathbf{M} = \{\mathbf{C}\} \cup S$, in which \mathbf{C} is a clause and S is a set of clauses. Note that it is perfectly fine for S to be empty. In the same way we can visualize $\mathbf{C} = \{s \approx t\} \cup \mathcal{E}$ by using vector column notation. By combining these we can pattern match against certain *literals* and sub-clauses/matrices that are of interest. For example

$$\mathbf{M} = \left[\begin{array}{c} [s \approx t] \\ \mathcal{E} \end{array} \right] S$$

lets us match a clause (\mathbf{C}) consisting of the literal $s \approx t$, the remaining (possibly empty) literal set \mathcal{E} and the remaining (possibly empty) set of clauses S .

3.1.2. Most General Unifiers

Given a set of equations $E = \{s_1 \approx t_1, \dots, s_n \approx t_n\}$, a unifier of E is a (variable) substitution σ such that $\sigma(s_1) = \sigma(t_1), \dots, \sigma(s_n) = \sigma(t_n)$. Given the set of unifiers $U(E)$ of E , a substitution σ is a *most general unifier (mgu)* of E if for all $\sigma' \in U(E)$, σ' is an instance of σ .

3.1.3. Rules

The following preprocessing rules should be read top down. If the conditions presented above the bar hold then one can infer the result below. $\mathbf{M}_1 \models \mathbf{M}_2$ means that \mathbf{M}_2 is a logical consequence of \mathbf{M}_1 , $\models \mathbf{M}$ means that \mathbf{M} is valid.

3.2. Valid Clauses

Any matrix that contains a clause \mathbf{C} consisting only of positive equalities is valid if there exists a most general unifier (mgu) for \mathbf{C} . Note that by this rule, an empty clause is valid.

$$\frac{[\mathbf{C} \ S] \quad \exists \sigma : \sigma = mgu(\mathbf{C}) \quad \forall L \in \mathbf{C} : L = (s \approx t)}{\models [\mathbf{C}] \cup \alpha([\mathbf{C}])}$$

As all variables local to \mathbf{C} are existentially quantified an mgu for such a clause represents an assignment for every variable in \mathbf{C} such that every equality literal is true. If the mgu is empty then every element of \mathbf{C} must have been of the form $s_1 \approx s_1, \dots, s_n \approx s_n$. One can construct a

new matrix \mathbf{M}' consisting of the clause \mathbf{C} extended with the axioms of equality such that the resulting matrix is valid. As \mathbf{M}' is a subset of \mathbf{M} then $\mathbf{M} \cup \alpha(\mathbf{M})$ is valid. This rule not only provides a termination case for the reduction algorithm, but allows certain (artificial) theorems to be proven in a very efficient manner.

In order to ensure that the reader is under no illusion, we will take some time to clarify this first rule. There are three conditions that must be met, read from left to right they are: The presence of a clause \mathbf{C} . The existence of a substitution σ such that σ is the mgu for clause \mathbf{C} (this does not preclude the empty substitution $\sigma = \emptyset$). Finally, if the clause \mathbf{C} is not empty, then it must only contain positive equality literals. If these conditions are met, then we can construct a new matrix consisting of the single clause \mathbf{C} . If we union this matrix with the set of all axioms that \mathbf{C} generates then the result can be shown to be valid using first-order logic alone.

As an example, consider the following: $\{x \approx a, f(x) \approx f(a), \dots\}$. While there may be multiple ways to prove the validity of such a formula - we can see that all we need to show is that there exists a value x that is equal to a . Due to the reflexivity of equality such a search is rather straightforward. Even though the occurrence of such a formula may be uncommon, detecting such a clause can be crucial for a successful proof search. The exclusion of this rule can lead to a proof procedure timing out before *reaching* the clause in question.

3.3. Contradictions

When considering a formula that is equality free then any clause \mathbf{C} that contains $P(s)$ and $\neg P(s)$ is contradictory and can be removed. This idea can be extended to clauses containing $P(s)$ and $\neg P(t)$ if the remaining clause \mathcal{E} implies that s and t are equal. If we can derive $s \approx t$ in \mathcal{E} then both $(P(s) \wedge \neg P(t) \wedge \mathcal{E})$ and $(s \neq t \wedge \mathcal{E})$ must result in contradictory clauses.

$$\frac{\left[\begin{array}{c} P(s_1, \dots, s_n) \\ \neg P(t_1, \dots, t_n) \\ \mathcal{E} \end{array} \right] S \quad \mathcal{E} \models s_1 \approx t_1 \wedge \dots \wedge s_n \approx t_n}{\models S \cup \alpha(S) \quad \text{iff} \quad \models \mathbf{M} \cup \alpha(\mathbf{M})}$$

In the same way as for predicates, if we derive that $a \approx b$ in the remainder of the clause, then it can not be the case that both $a \approx b$ and $a \neq b$ can be true simultaneously.

$$\frac{\left[\begin{array}{c} s \neq t \\ \mathcal{E} \end{array} \right] S \quad \mathcal{E} \models s \approx t}{\models S \cup \alpha(S) \quad \text{iff} \quad \models \mathbf{M} \cup \alpha(\mathbf{M})}$$

3.4. Redundancy

Consider a clause that contains two equalities $a \approx b$ and $f(a) \approx f(b)$. If $a \approx b$ is true then $f(a) \approx f(b)$ follows, meaning that $f(a) \approx f(b)$ is redundant. However, $a \approx b$ does not follow from $f(a) \approx f(b)$. The same principle can be used to find *redundant* predicates.

$$\frac{\left[\begin{array}{c} s \approx t \\ \mathcal{E} \end{array} \right] S \quad \mathcal{E} \models s \approx t}{\models [\mathcal{E} \ S] \cup \alpha([\mathcal{E} \ S]) \quad \text{iff} \quad \models \mathbf{M} \cup \alpha(\mathbf{M})}$$

The rule of redundancy is an interesting one. While it may seem advantageous to minimise the number of literals in a clause, we must remember that our aim is to reduce the search space of the problem. It may well be the case that an equation or predicate that is redundant in terms of derivability, is key to the proof of a formula. If this is the case then the responsibility is on the theorem prover to - in some sense - re-find this literal.

3.5. Pure Clauses

Assume that $P(s) \in C_j$. If there does not exist a k such that $\neg P(t) \in C_k$ then C_j is isolated; the same holds for $\neg P(s) \in C_j$ if there is no $P(t)$. An isolated clause may be removed from the matrix. This general rule is also applicable to formulae that contain equality.

3.6. Unsatisfiable Clauses

If a matrix without equality axioms does not contain negated equality then the only instance (of negated equality) can come from the addition of the equality axioms. A negated equality has to be part of the connection to make an appropriate path complementary. As every equality axiom apart from reflexivity contains at least one positive equality again, the only possible connection *to make this path complementary* has to be to the literal in the reflexivity axiom. This is equivalent to finding the mgu for the positive equation. If no mgu exists then the clause is contradictory and can be removed.

$$\frac{\left[\begin{array}{c} s \approx t \\ \mathcal{E} \end{array} \right] S}{\models [\sigma(\mathcal{E}) \ S] \cup \alpha([\sigma(\mathcal{E}) \ S]) \quad \text{iff} \quad \models \mathbf{M} \cup \alpha(\mathbf{M})} \quad \exists \sigma : \sigma = mgu(s \approx t) \quad \forall (u, v) : (u \neq v \notin \mathbf{M})$$

$$\frac{\left[\begin{array}{c} s \approx t \\ \mathcal{E} \end{array} \right] S}{\models S \cup \alpha(S) \quad \text{iff} \quad \models \mathbf{M} \cup \alpha(\mathbf{M})} \quad \nexists \sigma : \sigma = mgu(s \approx t) \quad \forall (u, v) : (u \neq v \notin \mathbf{M})$$

Thus, if there exists a matrix \mathbf{M} that contains positive equality but does not contain negated equality, then it can be reduced to a matrix \mathbf{M}' that is equality free such that the transformation is sound and complete, i.e. validity preserving.

3.7. Unit Clause

The final rule that we will consider concerns negated unit clauses. Consider a formula of the form where s and t are variable-free:

$$(s \approx t) \Rightarrow P(s) \vee \neg P(t)$$

When converted to disjunctive normal form this will result in a matrix that contains a negated unit clause:

$$[[s \neq t] \ [P(s)] \ [\neg P(t)]]$$

If the negated equality clause is used in a proof then every path must pass through it.

If $s \neq t$ is contradictory then the original formula was of the form $(s \approx s) \Rightarrow F$, in which case the matrix becomes F . If this is not the case then the following holds

$$\frac{[[s \neq t] \quad S] \quad \text{vars}(s \neq t) = \emptyset}{\models \mathbf{M} \cup \alpha(\mathbf{M}) \quad \text{if} \quad \models S\{s \mapsto t\} \cup \alpha(S\{s \mapsto t\})}$$

$$\frac{[[s \neq t] \quad S] \quad \text{vars}(s \neq t) = \emptyset}{\models \mathbf{M} \cup \alpha(\mathbf{M}) \quad \text{if} \quad \models S\{t \mapsto s\} \cup \alpha(S\{t \mapsto s\})}$$

where $\text{vars}(s \neq t)$ is the set of all variables in $s \neq t$ and $S\{t \mapsto s\}$ is the set (of remaining clauses) S in which all terms t have been replaced by the term s . This rule is sound as every path through \mathbf{M} contains the negated equation $s \neq t$, and applications of the equality axioms can be used to replace any occurrence of s by t or vice versa. While soundness is preserved, completeness of the calculus is lost. While this may sound problematic, we will see that in “real world” situations such a limitation has little to no negative impact.

One special case is that when $s \neq s$ we know that $\models M \cup \alpha(\mathbf{M})$ iff $\models S \cup \alpha(S)$ by the rule of contradictions 3.3. Out of all of the rules discussed, this one is the most powerful. As we will see in the evaluation section, the ability to make *good* choices in terms of choosing a direction (or not applying the rule) is an important factor.

4. Implementation

The preprocessing rules of the previous section have been implemented in the EPICC (Equality Preprocessing in Connection Calculi) system.^{2 3} The aim of this work is to improve the performance of equality handling for connection calculi – with a particular focus on the leanCoP theorem prover. As the rules discussed in Section 3 will be run on large matrices they needed to be performant. Such an approach has certainly influenced both the design of the EPICC run-time system, and the considerations made when searching for rules to implement. The rules that we have seen in Section 3 were influenced not only by the nature of the connection calculus, but also the internal representation of the matrix and the fact that we plan on using it as a preprocessing step.

The current version of EPICC is written in the functional Lisp-like language Clojure that runs on the JVM (Java Virtual Machine). The data driven approach of the language makes it ideal for prototyping and exploration. The current rule based approach has been developed with portability in mind (an implementation in Haskell is being actively developed).

The current approach taken by the leanCoP theorem prover is as follows: If a matrix \mathbf{M} contains an equality symbol then the matrix is extended with the axioms of equality before the actual proof search starts, proving the matrix $\mathbf{M} \cup \alpha(\mathbf{M})$. EPICC replaces this procedure by first applying the rules discussed in Section 3, before generating the set of axioms $\alpha(\mathbf{M})$. For the current implementation, formulae in a disjunctive normal form matrix format are accepted.

²Available for download under the GPL license at <http://leancoP.de/epicc/>.

³The Clojure implementation of the preprocessing steps that does not include the leanCoP core prover can be obtained at <https://github.com/beoliver/clj-epicc>.

$$\frac{[C \ S] \quad \exists \sigma : \sigma = mgu(C) \quad \forall L \in C : L = (s \approx t)}{\models [C] \cup \alpha([C])}$$

```
(defrecord Valid-Clause []
  LocalRule
  (candidate-clauses [rule m]
    (filter #(only-pos-eq? (lookup m %)) (pos-eq-clauses m)))
  (apply-local-rule [rule m c]
    (cond (apply mgu? (formulae c)) (valid "...")
          (empty? (neg-eq-clauses m)) (redundant "...")
          :else (no-op)))
```

Figure 2: An example of how the rule for valid clauses can be encoded in the Clojure language.

Reductions are expressed in terms of *rules*. The concepts of programming against interfaces and using well-defined return values are common features of many popular languages (both functional and imperative). A local rule is something that implements two functions.

```
(defprotocol LocalRule
  (candidate-clauses [rule matrix])
  (apply-local [rule matrix clause]))
```

The same idea can be expressed in Haskell using data types.

```
data Rule = Rule { candidate_clauses :: Matrix -> [Int]
                  , apply_local :: Matrix -> Clause -> Result }
```

The function `candidate-clauses` is responsible for returning a sequence of all of the indexes that the rule may be applicable to. It is assumed that the search for candidates is cheap. By separating a rule in this way we gain the ability to arbitrarily terminate the reduction process (either due to time constraints or having found a solution) while retaining the most current version of the matrix. Moreover, rules can be added, deleted, and re-ordered in a straightforward manner aiding the development of new approaches.

As a concrete example, let us consider the rule for valid clauses and how it could be implemented in the Clojure language as shown in Figure 2.

The local rules are most commonly responsible for *deciding* if a given clause is valid, redundant, or contradictory. Deletion and termination is handled by the reduction function. Variants exist that allow a rule to return a new clause, for example removing some $s \approx s$. In this case the reduction function will update the matrix to reflect the changes.

4.1. The Supervisor process

A *supervisor* is used that manages both state changes and termination conditions in order to run the rules. This *supervisor* is implemented using a function that continually loops (Clojure’s “loop” keyword can be thought of as a recursive while loop that allows arguments to be passed) waiting for a termination condition. This means that a *local* rule is responsible for *deciding* if a

given clause is valid, redundant or contradictory. Any deletion, termination or global re-writing is handled by the *supervisor* not the rule.

The *supervisor* tracks the state of the matrix as well as keeping a history of all rules that were applicable (i.e. they did not return NoOp). One result of this is that if a certain application of rules yields a valid result, then the system can extract the submatrix from the original input, add the axioms of equality and pass it on to a theorem prover. The rules are run until a terminating condition is met. Such a condition may either be a timeout, reaching a certain number of iterations, finding a solution or the previous iteration producing no change. The axioms of equality are added to the resulting matrix when no more reductions can be performed. It is possible to disable the equality axiom generation if desired.

Because it is possible for a rule to return “valid” or “invalid”, the supervisor can be seen as a partial proof procedure. While this functionality is currently not used explicitly by EPICC it was noted that during testing, the procedure was able to prove the validity of a handful of problems from the TPTP library directly. In terms of current implementation, the resulting matrix is always passed to the leanCoP theorem prover. In the case of the rules alone deriving “valid”, the resulting matrix can be proven by leanCoP .

While this does not affect the correctness of the testing results, it does mean that the EPICC framework requires users to know about which rules they are planning on using. One would imagine that in a future version of EPICC this information would be handled by the supervisor. Not only would this allow for a minor optimization in not having to invoke a theorem prover in all cases, but it would make the system more user friendly.

4.2. Internal representations and data structures

Internally, clauses are indexed sets of literals. A matrix is represented as a mapping from clause indexes to the clauses as well as additional information such as the indexes of clauses that contain positive/negated equality and a mapping from terms to the clauses that they occur in (an important factor when it comes to performance). These tables are updated every time a clause is added/removed from the matrix, meaning that when implementing a rule that, e.g., only considers clauses containing positive equality, one does not need to test every clause. The same principle is true when performing global rewriting – one only needs to update the clauses that a term occurs in.

When the matrix is imported, addition information is gathered, such as indexes of clauses that contain positive/negated equality. An internal table is also built up mapping terms to the clauses that they occur in. The reason for this table is to improve performance when performing term substitutions – by knowing which clauses contain a term s , the cost of applying the substitution $\{s \rightarrow t\}$ is reduced to the number of clauses that currently contain the term s (as opposed to naïvely applying the substitution to every clause). These tables are updated every time a clause is added/removed from the matrix, meaning that when implementing a rule that, say, only considers clauses containing positive equality, one does not need to repeatedly check every clause in the input matrix.

5. Evaluation

The equality preprocessing techniques described in Section 3 and implemented in Section 4 were evaluated on all relevant 8044 first-order (so-called FOF) formulae or problems contained in the TPTP library v6.4.0.

If a problem did not contain equality, or it caused a parser error then it was ignored. This resulted in a set of 4672 problems that would be used for testing. For every problem we have verified that the results, i.e. *theorem* or *non-theorem*, are consistent with the TPTP status of the problem. Of the 4672 problems that the tests were run on, we are particularly interested in the 4189 that have the TPTP status of *Theorem*.

5.1. Method

Six transformation strategies were compared against each other. The Clojure implementation of the EPICC system was used to perform all strategies. The first of these strategies simply adds the axioms of equality. This is the current approach taken by leanCoP and as such the results would provide a baseline for the remaining five approaches. The second approach performed the modification method (STE). As this is a well known approach for handling equality that can be performed during preprocessing it is interesting to see how it compares to the EPICC techniques described in this paper. The third approach would not perform any translation, nor would it add the axioms of equality. The remaining three configurations were variations of the EPICC techniques. Such a selection provides us with the ability to both compare the transformations discussed with the current leanCoP technique and to compare the transformations with two more approaches.

Each of the six methods were evaluated on all relevant problems contained in the TPTP library v6.4.0 [1]. TPTP is a library of test problems which supports the testing and evaluation of ATP systems. The library is divided into problem “domains”. Some examples of the domains are software verification (SWV and SWW), where it is formally established that a computer program does the task it is designed for, software creation (SWC), which is used to form a computer program that meets given specifications, general algebra (ALG), category theory (CAT), geometry (GEO), graph theory (GRA), knowledge representation (KRS), management (MGT), number theory (NUM), puzzles (PUZ), ring theory (RNG), set theory (SET and SEU), and syntactic (SYN).

As opposed to reading problems directly from the TPTP library, leanCoP was used to first convert each of the problems into disjunctive normal form. If a formula did not contain equality then it would not be used in the benchmarking tests. The resulting matrices were saved to disk and individually read by EPICC.

Every problem in the TPTP library has a status. We will concern ourselves with “Theorem” and “Non-Theorem”. For every problem it was recorded whether the result produced by a method is coherent with the TPTP status of the problem. This allows us to see which methods result in errors due to the method not preserving the completeness of the input formula. Such an approach also allows us to verify that no method implementation results in an unsound theorem prover.

The six approaches to be used were:

- **Axioms AX.** These results are used to provide a set of baseline results that other methods can be compared against. The input file is to be read by EPICC and the axioms of equality added before passing the resulting matrix to leanCoP. This corresponds to the technique the full leanCoP prover uses.
- **Modification method MM.** The input is read and Brand’s modification method performed on the matrix before passing it to leanCoP. As we have seen, the modification method is an existing preprocessing technique that can be used to *eliminate* equality. An implementation was written that would accept matrices in disjunctive normal form. It should be noted that this implementation is naïve in the sense that no attempt was made to optimize it. The algorithm is based on the one outlined in Brand’s paper [14].
- **No axioms NO-AX.** The file is read by EPICC and the matrix passed to leanCoP without adding any axioms of equality. The reason for including this approach is to provide an insight into the practical need for explicit equality handling. If a formula can be shown to be a theorem without adding the axioms of equality (i.e. without interpretation of equality) then the clauses containing equality were redundant. This approach cannot guarantee to preserve the completeness for formulae containing equality.
- **EPICC-1.** A configuration of EPICC that preserves the completeness for formulae containing equality. This is achieved by not applying the *Unit Clause* rule. The following (complete) rules are used: *Valid Clauses*, *Contradictions*, *Pure Clauses*, and *Unsatisfiable Clauses*. The axioms of equality are added after no more transformations can be applied.
- **EPICC-2.** A configuration of EPICC that uses a left-to-right rewrite rule for *Unit Clause*. The following rules were used: *Valid Clauses*, *Contradictions*, *Pure Clauses*, *Unsatisfiable Clauses*, and the *Unit Clause* rule, which performs rewriting in a left-to-right manner meaning that $(s \approx t)$ would result in a global substitution $\sigma = \{s \mapsto t\}$. The axioms of equality are added after no more transformations can be applied.
- **EPICC-3.** A configuration of EPICC that uses a custom rewrite rule for *Unit Clause*. The axioms of equality are added after no more transformations can be applied. The difference between this approach and that of **EPICC-2** is how the *Unit Clause* rule decides if a clause would result in the substitution $\sigma = \{s \mapsto t\}$, the substitution $\sigma = \{t \mapsto s\}$, or if the unit clause should be left alone. For example, one of the conditions is that a substitution $\sigma = \{s \mapsto t\}$ is only allowed if s does not occur in t .

For every input problem, six different outputs were produced – each corresponding to one of the strategies listed above. The core (Prolog) prover of leanCoP 2.1 using its complete *core* strategy (i.e. “[cut, comp(7)]”, see [10]) was then to be invoked for each output and allowed to run for at most ten seconds before being cancelled. The used core strategy uses *restricted backtracking* [12], which is switched off when a proof search depth of seven is reached. If a timeout occurs, then “timeout” was recorded. The core prover does not add any axioms of equality. SWI-Prolog version 8.0.2 was used for running leanCoP.

The choice to avoid using the strategy scheduling features of leanCoP was made for two reasons. As the strategy [cut, comp(7)] is complete we know that if running leanCoP on the output of one of the six configurations being tested results in “Non-Theorem” and the TPTP

library has it marked with the status “Theorem”, then that configuration did not preserve the completeness of the input formula. If we were to use strategy scheduling then leanCoP would ignore the result “Non-Theorem” if the internal strategy that proved the result was not complete. While we would eventually achieve the same result due to the last strategy that leanCoP employs (i.e. “[def]”) being complete, we might run out of time before this occurred. Secondly, we are interested in the effect that the transformations themselves have on the performance. It is assumed that if a particular leanCoP strategy is effective then that improvement would be seen across all transformations. Such a decision is certainly open to debate. It may well be the case that a particular strategy of leanCoP amplifies the effect of a particular transformation. However, as we wish to generalise the implementation of the theorem prover, such an event is beyond the scope of this work.

All evaluations were conducted on a six-core 2.2 GHz Intel Core i7 Macbook Pro with 16 GB of RAM running MacOS 10.14.4. Each input file was parsed by the Clojure implementation of the EPICC system and any transformations performed before invoking leanCoP on the output. EPICC was run using Clojure 1.10.0 and Java 1.8.0_181. The CPU time limit for each proof attempt by the leanCoP (core) prover was set to 10 seconds. When calculating the amount of time that a proof attempt takes, the results do *not* take into account the amount of time that was spent performing the preprocessing by EPICC.

5.2. Experimental Results

Of the 8044 FOF problems contained in the TPTP library 4672 problems contain equality. Of the 4672 problems containing equality that the tests were run on, we are particularly interested in the 4189 problems that have the TPTP status of “Theorem”. We will consider all results with respect to these 4189 problems. The reason for only considering the results of formulae that have the TPTP status of “Theorem” is because three of the methods tested (**NO-AX**, **EPICC-2** and **EPICC-3**) are known to not preserve the completeness of the input formula in the general case. In the case of the transformation **NO-AX** it is simply because it does not include the axioms of equality. For the other two it is due to their use of the *Unit Clause* rule. Thus, if we get the result “Non-Theorem” when using these methods, we do not know if it was derived because the original formula in the TPTP library is a non-theorem, or if it was due to the not completeness preserving transformation method.

Table 1 provides an overview of the results. Table 2 presents the results aggregated by problem rating, while Table 3 presents the results broken down by domain. The *rating* of a problem describes how difficult it is to solve it. E.g., a rating of 0.0 means that the problem is solved by all state-of-the-art provers, a rating of 0.7 means that it is solved by 30% of them.

The three methods that use techniques described in this paper (**EPICC-1**, **EPICC-2** and **EPICC-3**) outperform all other methods including the standard leanCoP approach (**AX**) in the number of theorems proven under ten seconds. Of all the methods tested, the modification method (**MM**) has the worst performance. While no optimizations were made to this method, the fact that the standard approach of **AX** (that also includes no optimization) resulted in 795 more proofs certainly suggests leanCoP handles the increase of search space introduced by the axioms of **AX** better than the increase of search space due to the large number of new clauses introduced by **MM**.

Table 1

Summary of different equality handling techniques for the 4189 considered problems.

	AX	MM	NO-AX	EPICC-1	EPICC-2	EPICC-3
proved	969	174	725	973	1002	1099
0 to 1sec.	814	116	629	814	841	922
1 to 10sec.	154	58	96	159	161	177
not proven by AX	-	11	31	9	109	166

Table 2

All theorems proved under 10 seconds aggregated by problem rating. A **bold** font denotes the best performing method(s).

		AX	MM	NO-AX	EPICC-1	EPICC-2	EPICC-3
rating	total						
0.0 to 0.09	706	496	144	373	498	476	524
0.1 to 0.19	441	221	25	147	219	238	259
0.2 to 0.29	388	96	4	73	98	112	124
0.3 to 0.39	351	73	0	68	74	88	94
0.4 to 0.49	374	46	0	32	47	49	52
0.5 to 0.59	360	22	0	16	22	23	30
0.6 to 0.69	275	8	0	7	8	9	9
0.7 to 0.79	297	6	0	8	6	6	6
0.8 to 0.89	324	1	0	1	1	1	1
0.9 to 0.99	247	0	0	0	0	0	0
1.0	426	0	0	0	0	0	0
total	4189	969	173	725	973	1002	1099

Of the three EPICC approaches, the performance of the complete method **EPICC-1** is only marginally better than **AX**. This suggests that (for the problems considered) the re-writing rule *Unit Clause* has a major impact. It would seem that either the (complete) methods described in Section 3 do little to reduce search space, or that the conditions required for their application are unfortunately not met frequently. **EPICC-2** performs better than **EPICC-1**. **EPICC-2** uses a left-to-right rewrite strategy for the *Unit Clause* rule. The method that leads to the most proofs being found is **EPICC-3**. **EPICC-3** uses a slightly less aggressive rewrite strategy for the *Unit Clause* compared to **EPICC-2**.

Figure 3 presents a visualization of the 200 proofs found that could not be found using the standard **AX** approach. This graphic can be thought of as an alternative to a Venn diagram. Multiple dots underneath a bar indicate that those solutions were found by multiple methods. The *intersection* row shows the cardinality of the intersections. The *solutions* column indicates the number of proofs that a particular method found that the standard **AX** approach could not. For example - of the 166 solutions that **EPICC-3** found, 97 were found by exactly one other method, namely **EPICC-2**.

Table 3

All theorems proved under 10 seconds aggregated by problem domain. A **bold** font denotes the best performing method(s).

		AX	MM	NO-AX	EPICC-1	EPICC-2	EPICC-3
Domain	total						
AGT	52	17	9	17	17	17	17
ALG	148	25	2	12	25	18	25
CAT	13	2	0	2	2	2	2
COM	24	6	2	5	6	6	6
CSR	28	3	1	1	3	2	3
GEO	114	9	1	6	9	10	10
GRA	18	3	0	3	4	3	4
GRP	79	6	0	4	6	5	5
HAL	6	1	0	1	1	1	1
KLE	224	16	4	3	16	14	16
KRS	14	9	0	10	10	10	10
LAT	105	14	6	13	16	16	16
MGT	45	16	6	13	16	14	17
MSC	1	1	1	0	1	1	1
NLP	15	1	0	2	1	1	1
NUM	558	155	23	122	152	145	158
PRO	63	9	0	11	9	9	9
PUZ	10	3	3	0	3	1	3
REL	108	1	0	0	1	0	1
RNG	152	32	5	24	34	29	39
SCT	79	6	0	6	6	6	6
SET	441	194	37	147	194	179	197
SEU	750	195	22	131	195	183	196
SWB	113	14	5	8	14	14	14
SWC	422	32	11	1	32	116	136
SWV	318	154	30	148	155	155	161
SWW	248	33	1	31	33	33	33
SYN	14	6	4	0	6	6	6
TOP	27	6	0	4	6	6	6
total	4189	969	173	725	973	1002	1099

We can see that not only does **EPICC-3** find the most solutions, but that many of the alternative methods appear to be subsumed by it. Indeed **EPICC-3** found 54 solutions that no other method could. **EPICC-2** only managed to find a single (unique) solution that **EPICC-3** (or any other method) could not. **EPICC-1** found no unique solutions.

The modification method **MM** returned 6 unique solutions which is quite interesting considering that it only found 11 solutions in total that the axiomatic approach could not. Indeed its general performance was poor, managing to find only 174 proofs (roughly 16% of the total achieved by **EPICC-3**) and never managing to find a solution for a problem with a rating higher than 0.29. However, it found new proofs that no other method could produce.

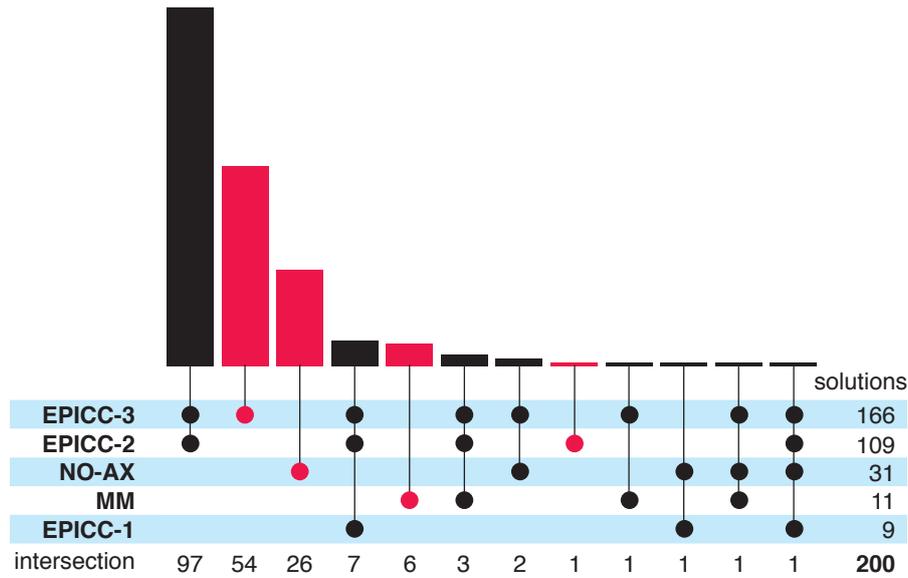


Figure 3: Visualising the “uniqueness” of the 200 proofs that could not be found using the standard leanCoP approach (**AX**).

Table 4

Evaluation of false negatives returned for the 4189 considered problems. A false negative occurs when an input theorem is transformed into an invalid formula.

	AX	MM	NO-AX	EPICC-1	EPICC-2	EPICC-3
false negatives	0	0	508	0	2	0
% of results	0%	0%	12%	0%	0.05%	0%

The results of the **NO-AX** approach are worthy of note. This method managed to prove a total of 31 theorems that the standard approach could not, with 26 of those being unique to this method alone. Such a high percentage (84%) of unique solutions is not that surprising if we consider the fact that by not adding axioms the search space is drastically reduced. The fact that **NO-AX** returned 508 false negatives (Table 4) supports this assumption. Indeed it is worthy of note that the **NO-AX** approach proved two more theorems within the 0.7 rating range than any other method, namely problems SEU205+1 (rating 0.77) and SEU241+2 (rating 0.73).

We see that for the most part the proofs of **EPICC-3** are a superset of the proofs of **EPICC-2** and **EPICC-1**. Thus a combination of **EPICC-3**, **NO-AX** and **MM** would find 199 of the 200 solutions. While the **EPICC-3** approach either performed equally as well as or (at least slightly) outperformed the other methods across all domains (Table 3), it was in the domain of Software Creation (SWC) that its performance was notably strong. This is shown in Figure 4, which shows the percentage of theorems proven with respect to the problem rating over all domains (left graph) and over the SWC domain only (right graph).

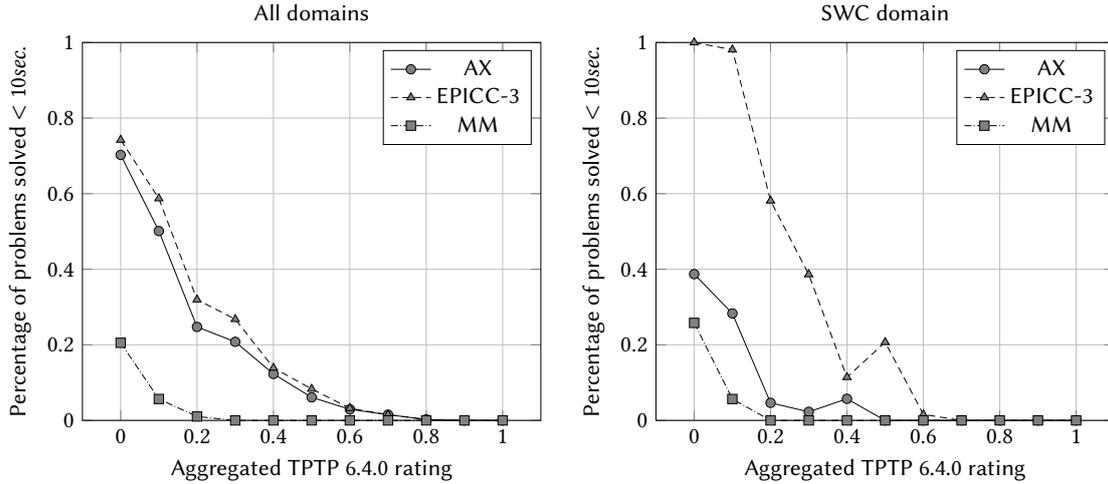


Figure 4: Percentage of theorems proven w.r.t. problem rating.

6. Conclusion

The present paper introduces EPICCC, a preprocessing technique for dealing with equality when proving formulae in (classical) first-order logic. Even though this technique can be used with any proof search calculus, it is in particular useful for tableau or connection calculi as the integration of techniques for equality into these calculi is not straightforward.

The preprocessing technique has been specified using a set of rules for simplifying or modifying a matrix representing the original formula in clausal form. The rules have been implemented, tested with the connection prover leanCoP and compared to the modification method and the standard approach of adding the equality axioms to the original matrix.

Using the EPICCC approach, leanCoP was able to prove significantly more problems of the TPTP library than using its standard technique of just adding the equality axioms, in particular in the “Software Creation” (SWC) domain. An interesting, yet to our knowledge so far undocumented, fact is that many of the problems in the TPTP library containing equality can be proved without any equality handling, i.e. treating the equality symbol as uninterpreted predicate symbol. The modification method proves significantly less problems than all other approaches. While the performance of the modification method may look undesirable, there were a handful of instances when it was the only approach that yielded a solution.

Future research work includes extending and optimizing the existent preprocessing rules. Furthermore, the adaptation and integration of similar preprocessing techniques into the non-clausal connection prover nanoCoP [17] or the non-classical provers ileanCoP and MleanCoP for first-order intuitionistic and modal logic [18, 19] is currently investigated.

Acknowledgments

We would like to thank the reviewers of a previous version of this paper for their comments.

References

- [1] G. Sutcliffe, The TPTP Problem Library and Associated Infrastructure. From CNF to TH0, TPTP v6.4.0, *Journal of Automated Reasoning* 59 (2017) 483–502.
- [2] G. Robinson, L. Wos, Paramodulation and theorem-proving in first-order theories with equality, in: J. H. Siekmann, G. Wrightson (Eds.), *Automation of Reasoning: 2: Classical Papers on Computational Logic 1967–1970*, Springer, Heidelberg, 1983, pp. 298–313.
- [3] J. A. Robinson, A machine-oriented logic based on the resolution principle, *Journal of ACM* 12 (1965) 23–41.
- [4] R. Hähnle, Tableaux and related methods, in: A. Robinson, A. Voronkov (Eds.), *Handbook of Automated Reasoning*, volume I, Elsevier Science, Amsterdam, 2001, pp. 101–178.
- [5] W. Bibel, Matings in matrices, *Commun. ACM* 26 (1983) 844–852.
- [6] J. Gallier, P. Narendran, S. Raatz, W. Snyder, Theorem proving using equational matings and rigid E-unification, *J. ACM* 39 (1992) 377–430.
- [7] A. Degtyarev, A. Voronkov, Simultaneous rigid E-unification is undecidable, in: H. Kleine Büning (Ed.), *Computer Science Logic*, Springer, Heidelberg, 1996, pp. 178–190.
- [8] P. Backeman, P. Rümmer, Efficient algorithms for bounded rigid E-unification, in: H. De Nivelle (Ed.), *Automated Reasoning with Analytic Tableaux and Related Methods*, Springer, Heidelberg, 2015, pp. 70–85.
- [9] P. Backeman, P. Rümmer, Theorem proving with bounded rigid E-unification, in: A. P. Felty, A. Middeldorp (Eds.), *Automated Deduction – CADE-25*, Springer, Heidelberg, 2015, pp. 572–587.
- [10] J. Otten, leanCoP 2.0 and ileanCoP 1.2: High performance lean theorem proving in classical and intuitionistic logic, in: A. Armando, P. Baumgartner, G. Dowek (Eds.), *IJCAR 2008*, volume 5195 of *LNAI*, Springer, Heidelberg, 2008, pp. 283–291.
- [11] J. Otten, W. Bibel, leanCoP: lean connection-based theorem proving, *Journal of Symbolic Computation* 36 (2003) 139–161.
- [12] J. Otten, Restricting backtracking in connection calculi, *AI Commun.* 23 (2010) 159–182.
- [13] G. Sutcliffe, The 9th IJCAR Automated Theorem Proving System Competition - CASC-29, *AI Communications* 31 (2018) 495–507.
- [14] D. Brand, Proving theorems with the modification method, *SIAM Journal on Computing* 4 (1975) 412–430.
- [15] W. Bibel, *Automated theorem proving*, Artificial intelligence, 2nd ed., F. Vieweg und Sohn, Wiesbaden, 1987.
- [16] P. B. Andrews, Refutations by matings, *IEEE Transactions on Computers* C-25 (1976) 801–807. doi:10.1109/TC.1976.1674698.
- [17] J. Otten, nanoCoP: A non-clausal connection prover, in: N. Olivetti, A. Tiwari (Eds.), *IJCAR 2016*, volume 9706 of *LNAI*, Springer, Heidelberg, 2016, pp. 302–312.
- [18] J. Otten, W. Bibel, Advances in connection-based automated theorem proving, in: J. Bowen, M. Hinchey, E.-R. Olderog (Eds.), *Provably Correct Systems*, NASA Monographs in Systems and Software Engineering, Springer, Heidelberg, 2017, pp. 211–241.
- [19] W. Bibel, J. Otten, From Schütte’s formal systems to modern automated deduction, in: R. Kahle, M. Rathjen (Eds.), *The Legacy of Kurt Schütte*, Springer, Heidelberg, 2020, pp. 217–251.