

REST: Integrating Term Rewriting with Program Verification

Anonymous author

Anonymous affiliation

Abstract

We introduce REST, a novel term rewriting technique for theorem proving that uses online termination checking and can be integrated with existing program verifiers. REST enables flexible but terminating term rewriting for theorem proving by: (1) exploiting newly-introduced term orderings that are more permissive than standard rewrite simplification orderings; (2) dynamically and iteratively selecting orderings based on the path of rewrites taken so far; and (3) integrating external oracles that allow steps that cannot be justified with rewrite rules. Our REST approach is designed around an easily implementable core algorithm, parameterizable by choices of term orderings and their implementations; in this way our approach can be easily integrated into existing tools. We implemented REST as a Haskell library and incorporated it into Liquid Haskell’s evaluation strategy, extending Liquid Haskell with rewriting rules. We evaluated our REST implementation by comparing it against both existing rewriting techniques and E-matching and by showing that it can be used to supplant manual lemma application in many existing Liquid Haskell proofs.

2012 ACM Subject Classification Theory of computation → Program verification

Keywords and phrases term rewriting, program verification, theorem proving

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

For all disjoint sets s_0 and s_1 , the identity $(s_0 \cup s_1) \cap s_0 = s_0$ can be proven in many ways. Informally accepting this property is easy, but a machine-checked formal proof may require the instantiation of multiple set theoretic axioms. Analogously, further proofs relying on this identity may themselves need to apply it as a previously-proven lemma. For example, proving functional correctness of any program that relies on a set data structure typically requires the instantiation of set-related lemmas. Manual instantiation of such universally quantified equalities is tedious, and the burden becomes substantial for more complex proofs: a proof author needs to identify exactly which equalities to instantiate and with which arguments; in the context of program verification, a wide variety of such lemmas are typically available. Given this need, most program verifiers provide some automated technique or heuristics for instantiating universally quantified equalities.

For the wide range of practical program verifiers that are built upon SMT solvers (e.g., [33, 23, 49, 37, 45, 43]), quantified equalities can naturally be expressed in the SMT solver’s logic. However, relying solely on such solvers’ E-matching techniques [19] for quantifier instantiation (as the majority of these verifiers do) can lead to both non-termination and incompletenesses that may be unpredictable [32] and challenging to diagnose [7]. The theory of how to prove that an E-matching-based encoding of equality reasoning guarantees termination and completeness is difficult and relatively unexplored [21].

A classical alternative approach to automating equality reasoning is *term rewriting* [26], which can be used to encode lemma properties as (directed) rewrite rules, matching terms against the existing set of rules to identify potential rewrites; the termination of these systems is a well-studied problem [16]. Although SMT solvers often perform rewriting as an internal simplification step, verifiers built on top typically cannot access or customize these rules, e.g., to add previously-proved lemmas as rewrite rules. By contrast, many



© Anonymous author(s);

licensed under Creative Commons License CC-BY 4.0



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

46 mainstream proof assistants (e.g., Coq [11], Isabelle/HOL [38], Lean [5]) provide automated,
 47 customizable term rewriting tactics. However, the rewriting functionalities of mainstream
 48 proof assistants either do not ensure the termination of rewriting (potentially resulting in
 49 divergence, for example Isabelle) or enforce termination checks that are overly restrictive in
 50 general, potentially rejecting necessary rewrite steps (for example, Lean).

51 In this paper, we present *REST* (*REwriting and Selecting Termination orderings*): a novel
 52 technique that equips program verifiers with automatic lemma application facilities via term
 53 rewriting, enabling equational reasoning with complementary strengths to E-matching-based
 54 techniques. While term rewriting in general does not guarantee termination, our tech-
 55 nique weaves together three key technical ingredients to automatically generate and explore
 56 guaranteed-terminating restrictions of a given rewriting system while typically retaining
 57 the rewrites needed in practice: (1) REST compares terms using well-quasi-orderings derived
 58 from (strict) simplification orderings; thereby facilitating common and important rules such
 59 as commutativity and associativity properties. (2) REST simultaneously considers an entire
 60 family of term orderings; selecting the appropriate term ordering to justify rewrite steps
 61 *during term rewriting itself*. (3) REST allows integration of an *external oracle* that generates
 62 additional steps outside of the term rewriting system. This allows the incorporation of rea-
 63 soning steps awkward or impossible to justify via rewriting rules, all without compromising
 64 the termination and relative completeness guarantees of our overall technique.

65 **Contributions and Overview** We make the following contributions:

- 66 1. We design and present a new approach (REST) for applying term rewriting rules and
 67 simultaneously selecting appropriate term orderings to permit as many rewriting steps
 68 as possible while guaranteeing termination (Sec. 3).
- 69 2. We introduce ordering constraint algebras, an abstraction for reasoning effectively about
 70 multiple (and possibly infinitely many) term orderings simultaneously (Sec. 4).
- 71 3. We introduce and formalize recursive path quasi-orderings (RPQOs) derived from the
 72 well-known recursive path ordering [15] (Sec. 4.1.2). RPQOs are more permissive than
 73 classical RPOs, and so let us prove more properties.
- 74 4. We formalize and prove key results for our technique: soundness, relative completeness,
 75 and termination (Sec. 5).
- 76 5. We implement REST as a stand-alone library, and integrate the REST library into Liquid
 77 Haskell to facilitate automatic lemma instantiation (Sec. 6).
- 78 6. We evaluate REST by comparing it to other term rewriting tactics and E-matching-based
 79 axiomatization, and show that it can substantially simplify equational reasoning proofs
 80 (Sec. 7).

81 We discuss related work in Sec. 8; we begin (Sec. 2) by identifying five key problems that
 82 all need solving for a reliable and automatic integration of term rewriting into a program
 83 verification tool.

84 **2 Five Challenges for Automating Term Rewriting**

85 In this section, we describe *five key challenges* that naturally arise when term rewriting is
 86 used for program verification and outline how REST is designed to address them. To illustrate
 87 the challenges, we use simple verification goals that involve uninterpreted functions and the
 88 set operators (\emptyset , \cup , \cap) that satisfy the standard properties of Figure 1. The variables x, y, z

<i>Name</i>	<i>Formula</i>
<i>idem-union</i>	$x \cup x = x$
<i>idem-inter</i>	$x \cap x = x$
<i>empty-union</i>	$x \cup \emptyset = x$
<i>empty-inter</i>	$x \cap \emptyset = \emptyset$
<i>commut-union</i>	$x \cup y = y \cup x$
<i>symm-inter</i>	$x \cap y = y \cap x$
<i>distrib-union</i>	$(x \cup y) \cap z = (x \cap z) \cup (y \cap z)$
<i>distrib-inter</i>	$(x \cap y) \cup z = (x \cup z) \cap (y \cup z)$
<i>assoc-union</i>	$x \cup (y \cup z) = (x \cup y) \cup z$

■ **Figure 1** Set identities used for examples in this section. Variables x, y, z are implicitly quantified. We write the binary functions \cup, \cap infix; along with (nullary) \emptyset these are fixed function symbols.

89 are implicitly quantified¹ in these rules. In formalizations of set theory, such properties may
 90 be assumed as (quantified) axioms, or proven as lemmas and then used in future proofs.

91 Term rewriting systems (defined formally in Sec. 5.1) are a standard approach for formally
 92 expressing and applying equational reasoning (rewriting terms via known identities). A term
 93 rewriting system consists of a finite set of *rewrite rules*, each consisting of a pair of a *source*
 94 *term* and a *target term*, representing that terms matching a rule's source can be replaced
 95 by corresponding terms matching its target. For example, the rewrite rule $x \cup \emptyset \rightarrow x$ can
 96 replace set unions of some set x and the empty set with the corresponding set x . Rewrite
 97 rules are applied to a term t by identifying some subterm of t which is equal to a rule's
 98 source under some substitution of the source's free variables (here, x , but not constants
 99 such as \emptyset); the subterm is then replaced with the correspondingly substituted target term.
 100 This rewriting step *induces an equality* between the original and new terms. For instance,
 101 the example rewrite rule above can be used to rewrite a term $f(s_0 \cup \emptyset)$ into $f(s_0)$, inducing
 102 an equality between the two.

103 Rewrite rules classically come with two restrictions: the free variables of the target
 104 must all occur in the source and the source must not be a single variable. This precludes
 105 rewrite rules which invent terms, such as $\emptyset \rightarrow x \cap \emptyset$, and those that trivially lead to infinite
 106 derivations. Under these restrictions, the first four identities induce rewrite rules from left-
 107 to-right (which we denote by e.g., *idem-inter* \rightarrow), while the remaining induce rewrite rules
 108 in both directions (e.g., *assoc-union* \rightarrow vs. *assoc-union* \leftarrow).

109 Next, we present a simple proof obligation taken from [34] in the style of equational
 110 reasoning (*calculational proofs*) supported in the Dafny program verifier [33].

111 ► **Example 1.** We aim to prove, for two sets s_0 and s_1 and some unary function f on sets,
 112 that, if the sets are disjoint (that is, $s_1 \cap s_0 = \emptyset$), then $f((s_0 \cup s_1) \cap s_0) = f(s_0)$.

$$\begin{aligned}
 \text{Equational Proof: } f((s_0 \cup s_1) \cap s_0) &= f((s_0 \cap s_0) \cup (s_1 \cap s_0)) && (\text{distrib-union}\rightarrow) \\
 &= f(s_0 \cup (s_1 \cap s_0)) && (\text{idem-inter}\rightarrow) \\
 &= f(s_0 \cup \emptyset) && (\text{disjointness ass.}\rightarrow) \\
 &= f(s_0) && (\text{empty-union}\rightarrow)
 \end{aligned}$$

(Possible Term Ordering, as explained shortly: RPO instance with $\cap > \cup$)

¹ over sets; we omit explicit types in such formulas, whose type-checking is standard.

XX:4 REST: Integrating Term Rewriting with Program Verification

114 This manual proof closely follows the user annotations employed in the corresponding
115 Dafny proof [34]; the application of the function f serves only to illustrate equational rea-
116 soning on subterms. Every step of the proof could be explained by term rewriting, hinting
117 at the possibility of an *automated* proof in which term rewriting is used to solve such proof
118 obligations. In particular, taking the term rewriting system naturally induced by the set
119 identities of Figure 1 *along with* the assumed equality expressing disjointness of s_0 and s_1
120 results in a term rewriting system in which the four proof steps are all valid rewriting steps.

121 In the remainder of the section, we consider what it would take to make term rewriting
122 effective for reliably automating such verification tasks. Perhaps unsurprisingly, there are
123 multiple problems with the simplistic approach outlined so far. The first and most serious
124 is that term rewriting systems in general *do not guarantee termination*; a proof search
125 may continue indefinitely by repeatedly applying rewrite rules. For example, the rules
126 *distrib-union* and *distrib-inter* can lead to an infinite derivation $(s_0 \cup s_1) \cap s_2 \rightarrow (s_0 \cap s_2) \cup$
127 $(s_1 \cap s_2) \rightarrow (s_0 \cup (s_1 \cap s_2)) \cap (s_2 \cup (s_1 \cap s_2)) \rightarrow \dots$

128 **Challenge 1:** Unrestricted term rewriting systems do not guarantee termination.

129 To ensure termination (as proved in Theorem 22) REST follows the classical approach of
130 restricting a term-rewriting system to a variant in which sequences of term rewrites (*rewrite*
131 *paths*) are allowed only if each consecutive pair of terms is *ordered* according to some term
132 ordering which rules out infinite paths.

133 For example, *Recursive path orderings* (RPOs) [15] define well-founded orders $>_{\mathcal{T}}$ on
134 terms \mathcal{T} based on an underlying well-founded strict partial order $>$ on *function symbols*.
135 Intuitively, such orderings use $>$ to order terms with different top-level function symbols,
136 combined with the properties of a *simplification order* [14] (e.g., compatibility with the
137 subterm relation). Different choices of the underlying $>$ parameter yield different RPO
138 instances that order different pairs of terms; in particular, potentially allowing or disallowing
139 certain rewrite paths.

140 In Example 1, the RPO based on the partial order $\cap > \cup$ and $\cap > \emptyset$ permits all the
141 rewriting steps, that is, the left-hand-side of each equation is greater than the right-hand-
142 side.

143 Sadly, this ordering will not permit the rewriting steps required by our next example.

144 ► **Example 2.** We aim to prove, for two sets s_0 and s_1 and some unary function f on sets,
145 that, if s_1 is a subset of s_0 (that is, $s_0 \cup s_1 = s_0$), then $f((s_0 \cap s_1) \cup s_0) = f(s_0)$.

$$\begin{aligned} \text{Equational Proof: } f((s_0 \cap s_1) \cup s_0) &= f((s_0 \cup s_0) \cap (s_1 \cup s_0)) && (\text{distrib-inter} \rightarrow) \\ &= f(s_0 \cap (s_1 \cup s_0)) && (\text{idem-union} \rightarrow) \\ &= f(s_0 \cap (s_0 \cup s_1)) && (\text{commut-union} \rightarrow) \\ &= f(s_0 \cap s_0) && (\text{subset ass.} \rightarrow) \\ &= f(s_0) && (\text{idem-inter} \rightarrow) \end{aligned}$$

(Possible Term Ordering: RPQO instance, explained shortly, with $\cap > \cup$)

147 An RPO based on the function symbol ordering $\cap > \cup$ (as required by Example 1) will
148 not permit the first step of this proof (since the RPO ordering first compares the top level
149 function symbols). Instead, this step requires an RPO based on the ordering $\cup > \cap$. To
150 accept *both* this proof step *and* the Example 1 we need different restrictions of the rewrite
151 rules for different proofs; in particular, different rewrite paths may be ordered according to
152 RPOs that are based on different function orderings.

153 To generalize this problem we will call RPOs a term ordering *family* that is *paramet-*
 154 *ric* with respect to the underlying function ordering. Thus, a concrete RPO term ordering
 155 (called an *instance* of the family) is obtained after the parametric function ordering is in-
 156 stantiated. With this terminology, the next challenge can be stated as follows:

Challenge 2: Different proofs require different term orderings within a family.

157
 158 Note that enumerating all term orderings in a term ordering family is typically impractical
 159 (this set is often very large and may be infinite). To address this challenge, REST uses a novel
 160 algebraic structure (Sec. 4.2) to allow for an abstract representation of sets of term orderings
 161 with which one can efficiently check whether any instance of a chosen term ordering family
 162 can orient the necessary rewrite steps to complete a proof.

163 Going back to Example 2, the RPO instance with $\cup > \cap$ will permit all the steps,
 164 apart from the commutativity axiom expressed by (*commut-union* \rightarrow). To permit this step
 165 we need an ordering for which $t_1 \cup t_2 >_{\mathcal{T}} t_2 \cup t_1$. But for RPO instances, as well as for
 166 many other term orderings, the terms $t_1 \cup t_2$ and $t_2 \cup t_1$ are equivalent and thus cannot be
 167 oriented; associativity axioms are also similarly challenging. Since many proofs require such
 168 properties, it is important in practice for rewriting to support them.

Challenge 3: Strict orderings restrict commutativity and associativity steps.

169
 170 To address this challenge REST relaxes the strictness constraint by requiring the chosen term
 171 ordering family to consist (only) of *thin well-quasi-orderings* (Sec. 4). Intuitively, such
 172 orderings permit rewriting to terms which are *equal* according to the ordering, but such
 173 equivalence classes of terms must be guaranteed to be finite. In Sec. 4 we show how to lift
 174 well-known families of term orderings to analogous and more-permissive families of thin well-
 175 quasi-orders. In particular, we show how to lift RPOs to a particularly powerful family of
 176 term orderings that we call *recursive path quasi-orderings (RPQOs)*, whose instances allow
 177 us to accept Example 2.

178 Despite the permissiveness of RPQOs, there remain some rewrite derivations that will
 179 be rejected by all term orderings in the RPQO family. For example, consider the following
 180 proof that set union is monotonic with respect to the subset relation:

181 ► **Example 3.** We aim to prove, for sets s_0 , s_1 , and s_2 , that, if s_1 is a subset of s_0 (that is,
 182 $s_0 \cup s_1 = s_0$), then $(s_2 \cup s_1) \cup (s_2 \cup s_0) = s_2 \cup s_0$.

$$\begin{aligned}
 \text{Equational Proof: } (s_2 \cup s_1) \cup (s_2 \cup s_0) &= s_2 \cup (s_1 \cup (s_2 \cup s_0)) && (\text{assoc-union}\leftarrow) \\
 &= s_2 \cup ((s_1 \cup s_2) \cup s_0) && (\text{assoc-union}\rightarrow) \\
 &= s_2 \cup ((s_2 \cup s_1) \cup s_0) && (\text{commut-union}\rightarrow) \\
 &= s_2 \cup (s_2 \cup (s_1 \cup s_0)) && (\text{assoc-union}\leftarrow) \\
 183 &= s_2 \cup (s_2 \cup (s_0 \cup s_1)) && (\text{commut-union}\rightarrow) \\
 &= s_2 \cup (s_2 \cup s_0) && (\text{subset ass.}\rightarrow) \\
 &= (s_2 \cup s_2) \cup s_0 && (\text{assoc-union}\rightarrow) \\
 &= s_2 \cup s_0 && (\text{idem-union}\rightarrow)
 \end{aligned}$$

(Possible Term Ordering: any KBQO instance)

184 The above rewrite rule steps cannot be oriented by any RPQO, but are trivially oriented
 185 by a quasi-ordering that is based on the syntactic size of the term, e.g., a quasi-ordering
 186 based on the well-known Knuth-Bendix family of term orderings [29]. Yet, a Knuth-Bendix
 187 quasi-ordering (KBQO, defined in Sec. 4) cannot be used on our previous two examples;
 188 fixing even a single choice of term ordering *family* would still be too restrictive in general.

Challenge 4: Some proofs require different families of term orderings.

189

190 To address this challenge, REST (Sec. 3.2) is defined parametrically in the choice and repre-
191 sentation of a term ordering family.

192 Finally, although equational reasoning is powerful enough for these examples, general
193 verification problems usually require reasoning beyond the scope of simple rewriting. For
194 example, simply altering Example 1 to express the disjointness hypothesis instead via car-
195 dinality as $|s_0 \cap s_1| = 0$ means that, to achieve a similar proof, reasoning within the theory
196 of sets is necessary to deduce that this hypothesis implies the equality needed for the proof;
197 this is beyond the abilities of term rewriting.

Challenge 5: Program verification needs proof steps not expressible by rewriting.

198

199 To address this challenge, our REST approach allows the integration of an external oracle that
200 can generate equalities not justifiable by term rewriting, while still guaranteeing termination
201 (Sec. 3.3).

3 The REST Approach

202

203 We develop REST to tackle the above five challenges and integrate a flexible, expressive,
204 and guaranteed-terminating term rewriting system with a verification tool. REST consists of
205 an interface for defining term orderings and an algorithm for exploring the rewrite paths
206 supported by the term orderings. In Sec. 3.1 we describe the representation of term orderings
207 in REST and how they address Challenges 2 and 4. In Sec. 3.2 we describe the REST algorithm
208 that is parametric to these orderings and Sec. 3.3 describes the integration with external
209 oracles (Challenge 5).

3.1 Representation of Term Orderings in REST

210

211 Rather than considering individual term orderings, REST operates on indexed sets (families)
212 of term orderings (whose instances must all be thin well-quasi-orderings).

213 ► **Definition 4** (Term Ordering Family). *A term ordering family Γ is a set of thin well-quasi-*
214 *orderings on terms, indexed by some parameters P . An instance of the family is a term*
215 *ordering obtained by a particular instantiation of P .*

216 For example, the concept of recursive path ordering is defined parametrically with respect
217 to a precedence on function symbols, and therefore defines a term ordering family indexed
218 by this choice of function symbol ordering.

219 A core concern of REST is determining whether any instance of a given term ordering
220 family can orient a rewrite path. However, term ordering families cannot directly compare
221 terms; doing so requires choosing an ordering inside the family. The root of Challenge 2
222 is that choosing an ordering in advance is too restrictive: different orderings are necessary
223 to complete different proofs. The idea behind REST's search algorithm is to address this
224 challenge by simultaneously considering all orderings in the family when considering rewrite
225 paths and continuing the path so long as it can be oriented by *any* ordering.

226 To demonstrate the technique, we show how REST's approach can be derived from a
227 naïve algorithm. The purpose of the algorithm is to determine if any ordering in a family Γ
228 can orient a path $t_1 \rightarrow \dots \rightarrow t_n$; i.e., if there is a $>_{\mathcal{T}} \in \Gamma$ such that $t_1 >_{\mathcal{T}} \dots >_{\mathcal{T}} t_n$.

<pre style="margin: 0;"> orients : (Set $O \times$ List \mathcal{T}) \rightarrow Bool orients(Γ, ts) = $os := \Gamma;$ for $i \in 1$ to $ts - 1$ { $os := \{>_{\mathcal{T}} \in os \mid ts_i >_{\mathcal{T}} ts_{i+1}\};$ if ($os = \emptyset$) return false; } return true; </pre>	(1)	<pre style="margin: 0;"> orients : (OCA \times List \mathcal{T}) \rightarrow Bool orients($(\top, refine, sat), ts$) = $c := \top;$ for $i \in 1$ to $ts - 1$ { $c := refine(c, ts_i, ts_{i+1});$ if (not($sat(c)$)) return false; } return true; </pre>	(2)	(3)
--	-----	---	-----	-----

■ **Figure 2** Two algorithms that determine if an ordering in the term ordering family Γ can orient a path of terms ts . **Left** presents the naïve, exhaustive algorithm. **Right** is using the ordering constraint algebra $\langle \top, refine, sat \rangle$ that returns true iff an ordering in Γ can orient ts without explicitly constructing any term orderings. Ois is the type of a term ordering.

229 The naïve algorithm is depicted on the left of Figure 2. The naïve algorithm works
 230 iteratively, computing the set of orderings os that can orient an increasingly-long path,
 231 short-circuiting if the set becomes empty. The algorithm enumerates each ordering in Γ
 232 and compares terms with each ordering (potentially multiple times). Unfortunately, this
 233 enumeration is not practical: some term ordering families have infinite or prohibitively large
 234 numbers of instances. REST avoids these issues by allowing the set of term orderings to be
 235 abstracted via a structure called an Ordering Constraint Algebra (OCA, Def. 14 of Sec. 4.2).

236 An OCA for a term ordering family Γ consists of a type C along with four parameters
 237 $\gamma : C \rightarrow \mathcal{P}(\Gamma)$, $\top : C$, $refine : C \rightarrow \mathcal{T} \rightarrow \mathcal{T} \rightarrow C$, and $sat : C \rightarrow Bool$. C is a type whose
 238 elements *represent* subsets of Γ . The function γ is the *concretisation function* of the OCA,
 239 not needed programmatically but instead defining the *meaning* of elements of C in terms
 240 of the subsets of the term ordering family they represent. The remaining three functions
 241 correspond to the operations on sets of term orderings used in lines (1), (2), and (3) of
 242 the naïve algorithm. \top represents the set of all term orderings in Γ , $refine(c, t, u)$ filters
 243 the set of orderings represented by c to include only those where $t >_{\mathcal{T}} u$, and $sat(c)$ is a
 244 predicate that returns true if the set of orderings represented by c is nonempty. Figure 2 on
 245 the right shows how the ordering constraint algebra can be used to perform an equivalent
 246 computation to the naïve algorithm, without explicitly instantiating sets of term orderings.
 247 The OCA plays a role similar to abstract interpretation in a program analysis, where C is
 248 an abstraction over sets of term orderings, and the results of the abstract operations on C
 249 correspond to their concrete equivalents. Namely, we have $\gamma(\top) = \Gamma$, $\gamma(refine(c, t_l, t_r)) =$
 250 $\{> \mid > \in \gamma(c) \wedge t_l > t_r\}$, and $sat(c) \Leftrightarrow \gamma(c) \neq \emptyset$.

251 The ordering constraint algebra enables three main advantages compared to direct com-
 252 putation with sets of term orderings:

- 253 1. The number of term orderings can be very large, or even infinite, thus making enumera-
 254 tion of the entire set intractable.
- 255 2. An OCA can provide efficient implementations for *refine* and *sat* by exploiting properties
 256 of the term ordering family. Comparing terms using the constituent term orderings
 257 requires repeating the comparison for each ordering, despite the fact that most orderings
 258 will differ in ways that are irrelevant for the comparison.
- 259 3. The OCA does not impose any requirements on the type of C or the implementation
 260 of \top , *refine*, and *sat*. For example, an OCA can use \top and *refine* to construct logical

<pre> REST : (OCA × R × T × (T → Set T)) → Set T REST(⟨T, refine, sat⟩, R, t₀, E) = o := ∅; p := [[t₀, T]]; while (p is not empty){ pop(ts, c) from p; t := last ts; o := o ∪ {t}; foreach (t' such that t' ∉ ts ∧ (t →_R t' ∨ t' ∈ E(t))){ if (t' ∈ E(t) ∨ (t →_R t' ∧ sat(refine(c, t, t')))){ push (ts ++ [t'], refine(c, t, t')) to p } } } return o; </pre>
--

■ **Figure 3** The REST algorithm.

261 formulas, with *sat* using an external solver to check their satisfiability. Alternatively,
 262 it could define *C* to be sets of term orderings that are reasoned about explicitly, and
 263 implement \top , *refine*, and *sat* as the operations of the naïve algorithm.


264 We now describe how the REST algorithm uses the OCA to explore rewrite paths.

265 3.2 The REST Algorithm

266 Figure 3 presents the REST algorithm. The algorithm takes four parameters. The first
 267 parameter is an OCA $\langle \top, \text{refine}, \text{sat} \rangle$, as discussed above. The algorithm's second parameter,
 268 *R*, is a finite set of term rewriting rules (not required to be terminating); for example, we
 269 could pass the oriented rewrite rules corresponding to Figure 1. The third parameter *t*₀ is
 270 the term from which term rewrites are sought. The final parameter *E* acts as an external
 271 oracle, generating additional rewrite steps that need *not* follow from the term rewriting rules
 272 *R*. To simplify the explanation, we will initially assume that $E = \lambda t. \emptyset$, i.e., this parameter
 273 has no effect. Our algorithm produces a set of terms, each of which are reachable by *some*
 274 rewrite path beginning from *t*₀, and for which *some* ordering allows the rewrite path. The
 275 algorithm addresses Challenge 1 (termination; Theorem 22) because every path must be
 276 finite: no ordering could orient an infinite path.

277 Our algorithm operates in worklist fashion, storing in *p* a list of pairs (*ts*, *c*) where *ts* is
 278 a non-empty list of terms representing a rewrite path already explored (the head of which
 279 is always *t*₀) and *c* tracks the ordering constraints of the path so far. The set *o* records the
 280 output terms (initially empty): all terms discovered (down any rewrite path) equal to *t*₀ via
 281 the rewriting paths explored.

282 While there are still rewrite paths to be extended, i.e., *p* is not empty, a tuple (*ts*, *c*) is
 283 popped from *p*. REST puts *t*, i.e., the last term of the path, into the set of output terms
 284 *o* and considers all terms *t'* that are: (a) not *already* in the path and (b) reachable by a
 285 single rewrite step of *R* (or returned by the function *E* explained later). The crucial decision
 286 of whether or not to extend a rewrite path with the additional step $t \rightarrow t'$ is handled in
 287 the if check of REST. This check is to guarantee termination, by enforcing that we only add



■ **Figure 4** A visualization of REST running on the term from Example 1. Each path through the tree shown represents a rewrite path uncovered by our algorithm; the edge labels show the rewrite rule applied. The red dotted lines indicate rewrite steps rejected by REST.

288 rewrite steps which would leave the extended path still justifiable by *some* term ordering,
 289 as enforced by the *sat* check.

290 Figure 4 visualizes the rewrite paths explored by our algorithm for a run correspond-
 291 ing to the problem from Example 1, using the OCA for the recursive path quasi-ordering
 292 (Sec. 4.2.1)². The manual proof in Example 1 corresponds to the right-most path in this tree;
 293 the other paths apply the same reasoning steps in different orders. In our implementation,
 294 we optimize the algorithm to avoid re-exploring the same term multiple times unless this
 295 could lead to further rewrites being discovered (cf. Sec. 6).

296 The arrow from the root of the tree to its child corresponds to the first rewrite REST
 297 applies: $f((s_0 \cup s_1) \cap s_0) \rightarrow f((s_0 \cap s_0) \cup (s_1 \cap s_0))$. This rewrite step can only be oriented
 298 by RPQOs with precedence $\cap > \cup$; therefore applying this rewrite constrains the set of
 299 RPQOs that REST must consider in subsequent applications. For example, the rewrite to the
 300 left child of $f((s_0 \cap s_0) \cup (s_1 \cap s_0))$ can only be oriented by RPQOs with precedence $\cup > \cap$.
 301 Since no RPQO can have both $\cap > \cup$ and $\cup > \cap$, no RPQO can orient the entire path from
 302 the root; REST must therefore reject the rewrite. On the other hand, the rewrite to the right
 303 child can be oriented by any RPQO where $s_0 > \emptyset$, $s_1 > \emptyset$, or $\cap > \emptyset$. The path from the root
 304 can thus continue down the right-hand side, as there are RPQOs that satisfy both $\cap > \cup$
 305 and the other conditions. The subsequent rewrites down the right-hand side do not impose
 306 any new constraints on the ordering: $f((s_0 \cap s_0) \cup \emptyset) >_{\mathcal{T}} f(s_0 \cap s_0) >_{\mathcal{T}} f(s_0)$ in all RPQOs.

307 Similarly, REST will prove Example 2 but will reject Example 3 when the input OCA
 308 represents RPQO orderings. As shown in our benchmarks (Table 2 of Sec. 7), Example 3 is
 309 solved by REST with an OCA for the Knuth-Bendix term ordering family.

310 3.3 Integrating an External Oracle

311 Finally, to tackle Challenge 5, we turn to the (so far ignored) third parameter of the algo-
 312 rithm, the external oracle \mathcal{E} . In the example variant presented at the end of Sec. 2, such
 313 a function might supply the rewrite step $s_0 \cap s_1 \rightarrow \emptyset$ by analysis of the logical assumption
 314 $|s_0 \cap s_1| = 0$, which goes beyond term-rewriting. More generally, any external solver capable
 315 of producing rewrite steps (equal terms) can be connected to our algorithm via \mathcal{E} . In our
 316 implementation in Liquid Haskell, we use the pre-existing *Proof by Logical Evaluation (PLE)*

² We omit the commutativity rules from this run, just to keep the diagram easy to visualize, but our implementation handles the example easily with or without them.

317 technique [50], which complements rewriting with the expansion of program function defini-
 318 tions, under certain checks made via SMT solving. Our only requirements on the oracle \mathcal{E}
 319 are that the binary relation on terms generated by calls to it is bounded (finitely-branching)
 320 and strongly normalizing (cf. Sec. 5).

321 Our algorithm therefore flexibly allows the interleaving of term rewriting steps and those
 322 justified by the external oracle; we avoid the potential for this interaction to cause non-
 323 termination by conditioning any further rewriting steps on the fact that the entire path
 324 (including the steps inserted by the oracle) can be oriented by at least one candidate term
 325 ordering.

326 The combination of our interfacing for defining term orderings via ordering constraint
 327 algebras, a search algorithm that effectively explores all rewrites enabled by the orderings,
 328 and the flexible possibility of combination with external solvers via the oracle parameter
 329 makes REST very adaptable and powerful in practice.

330 **4 Well-Quasi-Orderings and the Ordering Constraint Algebra**

331 Term orderings are typically defined as *strict well-founded* orderings; this requirement en-
 332 sures that rewriting will obtain a normal form. However, as mentioned in Challenge 3, the
 333 restriction to strict orderings limits what can be achieved with rewriting. In this section we
 334 describe the derivation of well-quasi-orderings from strict orderings (Sec. 4.1) and introduce
 335 Knuth-Bendix quasi-orderings (Sec. 4.1.1) and recursive path quasi-orderings (Sec. 4.1.2),
 336 two novel term ordering families respectively based on the classical recursive path and Knuth-
 337 Bendix orderings. In addition, we formally introduce ordering constraint algebras (Sec. 4.2)
 338 and use them to develop an efficient ordering constraint algebra for RPQOs.

339 **4.1 Well-Quasi-Orderings**

340 We define well-quasi-orderings in the standard way.

341 ► **Definition 5** (Well-Quasi-Orderings). *A relation \geq is a quasi-order if it is reflexive and*
 342 *transitive. Given elements t and u in S , we say $t \approx u$ if $t \geq u$ and $u \geq t$. A quasi-order \geq*
 343 *is also characterized as:*

- 344 1. *WQO, when for all infinite chains x_1, x_2, \dots there exists an $i, j, i < j$ such that $x_j \geq x_i$,*
- 345 2. *thin, when for all $t \in S$, the set $\{u \in S \mid t \approx u\}$ is finite, and*
- 346 3. *total, when for all $t, u \in S$ either $t \geq s$ or $s \geq t$.*

347 Well-quasi-orderings are not required to be antisymmetric, however the corresponding
 348 strict part of the ordering must be well-founded. Hence, a WQO derives a strict ordering
 349 over equivalence classes of terms; REST also requires that these equivalence classes are finite
 350 (i.e., the ordering is thin). With this requirement, REST guarantees termination by exploring
 351 only duplicate-free paths.

352 Many simplification orderings can be converted into more permissive WQOs. Intuitively,
 353 given an ordering $>_o$ its quasi-ordering derivation also accepts equal terms, so we denote it
 354 as \geq_o . We next present two such derivations.

355 **4.1.1 Knuth-Bendix Quasi-Orderings (KBQO)**

356 The Knuth-Bendix ordering [29] is a well-known simplification ordering used in the Knuth-
 357 Bendix completion procedure. Here, we present a simplified version of the ordering, used by
 358 REST that is using ordering to only compare ground terms.

359 ▶ **Definition 6.** A weight function w is a function $\mathcal{F} \rightarrow \mathbb{N}$, where $w(f) > 0$ for all nullary
 360 functions symbols, and $w(f) = 0$ for at most one unary function symbol. w is compatible with
 361 a quasi-ordering $\geq_{\mathcal{F}}$ on \mathcal{F} if, for any unary function f such that $w(f) > 0$, we have $f >_{\mathcal{F}} g$
 362 for all g . $w(t)$ denotes the weight of a term t , such that $w(f(t_1, \dots, t_n)) = w(f) + \sum_{1 \leq i \leq n} w(t_i)$

363 ▶ **Definition 7** (Knuth-Bendix ordering (KBO) on ground terms). The Knuth-Bendix Order-
 364 ing $>_{kbo}$ for a given weight function w and compatible precedence order $\geq_{\mathcal{F}}$ is defined as
 365 $f(t_1, \dots, t_m) = t >_{kbo} u = g(u_1, \dots, u_n)$ iff $w(t) \geq w(u)$, and:

- 366 1. $w(t) > w(u)$, or
- 367 2. $f >_{\mathcal{F}} g$, or
- 368 3. $f \geq_{\mathcal{F}} g$, and $(t_1, \dots, t_m) >_{kbolex} (u_1, \dots, u_n)$.

369 Where $>_{kbolex}$ performs a lexicographic comparison using $>_{kbo}$ as the underlying ordering.

370 Intuitively, KBO compares terms by their weights, using $\geq_{\mathcal{F}}$ and the lexicographic com-
 371 parison as “tie-breakers” for cases when terms have equal weights. However, as \geq is already
 372 a well-quasi-ordering on \mathbb{N} , we can derive a more general ordering by removing these tie-
 373 breakers and the need for a precedence ordering at all.

374 ▶ **Definition 8** (Knuth-Bendix Quasi-ordering (KBQO)). Given a weight function w , the
 375 Knuth-Bendix quasi-ordering \geq_{kbo} is defined as $t \geq_{kbo} u$ iff $w(t) \geq w(u)$.

376 The resulting quasi-ordering is considerably simpler to implement and is more permissive:
 377 $t >_{kbo} u$ implies $t \geq_{kbo} u$; and also enables arbitrary associativity and commutativity axioms
 378 as rewrite rules, since it only considers the weights of the function symbols and no structural
 379 components of the term. However, one caveat is that REST operates on well-quasi-ordering
 380 that are thin (Def. 5) and therefore can only consider KBQOs where $w(f) > 0$ for all unary
 381 function symbols f .

382 However, the fact that KBO and KBQO largely ignore the structure of the term in
 383 their comparison has a corresponding downside: it is not possible to orient distributivity
 384 axioms, or many other axioms that increase the number of symbols in a term. Therefore,
 385 we have found that a WQO derived from the recursive path ordering [15] to be more useful
 386 in practice.

387 4.1.2 Recursive Path Quasi-Orderings (RPQO)

388 In this section, we define a particular family of orderings designed to be typically useful for
 389 term-rewriting via REST. Our family of orderings is a novel extension of the classical notion
 390 of RPO, designed to also be more compatible with symmetrical rules such as commutativity
 391 and associativity (cf. Challenge 3, Sec. 2).

392 Like the classical RPO notions, our *recursive path quasi-ordering* (RPQO) is defined in
 393 three layers, derived from an underlying ordering on function symbols:

- 394 ■ The input ordering $\succ_{\mathcal{F}}$ can be any quasi-ordering over \mathcal{F} .
- 395 ■ The corresponding *multiset quasi-ordering* $\succ_{M(X)}$ lifts an ordering \succ_X over X to an
 396 ordering $\succ_{M(X)}$ over multisets of X . Intuitively $T \succ_{M(X)} U$ when U can be obtained
 397 from T by replacing zero or more elements in T with the same number of equal (with
 398 respect to \succ_X) elements, and replacing zero or more elements in T with a finite number
 399 of smaller ones (Def. 9).

XX:12 REST: Integrating Term Rewriting with Program Verification

400 ■ Finally, the corresponding *recursive path quasi-ordering* \succ_{rpo} is an ordering over terms.
 401 Intuitively $f(ts) \succ_{rpo} g(us)$ uses $\succ_{\mathcal{F}}$ to compare the function symbols f and g and the
 402 corresponding $\succ_{M(rpo)}$ to compare the argument sets ts and us (Def. 10).

403 Below we provide the formal definitions of the multiset quasi-ordering and recursive path
 404 quasi-ordering respectively generalized from the multiset ordering of [18] and the recursive
 405 path ordering [15] to operate on quasi-orderings. For all the three orderings, we write
 406 $x_l < x_r \doteq x_l \not\asymp x_r$ and $x_l > x_r \doteq x_l \succ x_r \wedge x_r \not\asymp x_l$.

407 ► **Definition 9** (Multiset Ordering). *Given a ordering \succ_X over a set X , the derived multiset*
 408 *ordering $\succ_{M(X)}$ over finite multisets of X is defined as $T \succ_{M(X)} U$ iff:*

- 409 1. $U = \emptyset$, or
- 410 2. $t \in T \wedge u \in U \wedge t \approx u \wedge (T - t) \succ_{M(X)} (U - u)$, or
- 411 3. $t \in T \wedge (T - t) \succ_{M(X)} (U \setminus \{u \in U \mid u <_X t\})$.

412 ► **Definition 10** (Recursive Path Quasi-Ordering). *Given a basic ordering $\succ_{\mathcal{F}}$, the recursive*
 413 *path quasi-ordering (RPQO) is the ordering \succ_{rpo} over \mathcal{T} defined as follows: $f(t_1, \dots, t_m) \succ_{rpo}$*
 414 *$g(u_1, \dots, u_n)$ iff*

- 415 1. $f >_{\mathcal{F}} g$ and $\{f(t_1, \dots, t_m)\} >_{M(rpo)} \{u_1, \dots, u_n\}$, or
- 416 2. $g >_{\mathcal{F}} f$ and $\{t_1, \dots, t_m\} \succ_{M(rpo)} \{g(u_1, \dots, u_n)\}$, or
- 417 3. $f \approx g$ and $\{t_1, \dots, t_m\} \succ_{M(rpo)} \{u_1, \dots, u_n\}$.

418 ► **Example 11.** As a first example, any RPQO $\succ_{\mathcal{T}}$ used to restrict term rewriting will
 419 accept the rule $x + y \rightarrow y + x$, since $x + y \succ_{\mathcal{T}} y + x$ always holds. Since the top level
 420 function symbol is the same $+ \approx +$, by Def. 10(3) we need to show $\{x, y\} \succ_{M(rpo)} \{y, x\}$.
 421 By Def. 9(2) (choosing both t and u to be x), we can reduce this to $\{y\} \succ_{M(rpo)} \{y\}$; the
 422 same step applied to y reduces this to showing $\emptyset \succ_{M(rpo)} \emptyset$ which follows directly from
 423 Def. 9(3).

424 From this example, we can see that both $x + y \succ_{rpo} y + x$ and $y + x \succ_{rpo} x + y$ hold, in
 425 this case independently of the choice of input ordering $\succ_{\mathcal{F}}$ on function symbols. In our next
 426 example, the choice of input ordering makes a difference.

427 ► **Example 12.** As a next example, we compare the terms $s(x)+y$ and $s(x+y)$. Now that the
 428 outer function symbols are *not* equal, the order relies on the ordering between $+$ and s . Let's
 429 assume that $+ >_{\mathcal{F}} s$. Now to get $s(x)+y \succ_{rpo} s(x+y)$, the 1st case of Definition 10 further
 430 requires $\{s(x)+y\} >_{M(rpo)} \{s(x+y)\}$, which holds if $s(x)+y >_{rpo} s(x+y)$. The outermost symbol
 431 for both expressions is $+$, so we must check the multiset ordering: $\{s(x), y\} >_{M(rpo)} \{s(x), y\}$,
 432 which holds because by case splitting on the relation between s and x , we can show that
 433 $s(x)$ is always smaller than x . In short, if $+ >_{\mathcal{F}} s$, then $s(x) + y \succ_{rpo} s(x + y)$.

434 Developing on our RPQO notion (Def. 10), we consider the set of *all* such orderings that
 435 are generated by any total, well-quasi-ordering over the operators. We prove that such term
 436 orderings satisfy the termination requirements of Theorem 22. Concretely:

437 ► **Theorem 13.** *If $\succ_{\mathcal{F}}$ is a total, well-quasi-ordering, then*

- 438 1. \succ_{rpo} is a well-quasi-ordering,
- 439 2. \succ_{rpo} is thin, and
- 440 3. \succ_{rpo} is thin well-founded.

441 **Proof.** The detailed proofs can be found in App. B. (1) uses the well-foundedness theorem
 442 of Dershowitz [15] and the fact that \succ_{rpo} is a quasi-simplification ordering. (2) relies on
 443 the fact that a finite number of function symbols can only generate a finite number of equal
 444 terms. (3) is a corollary of (1) and (2) combined. ◀

4.2 Ordering Constraint Algebras

Ordering constraint algebras play a crucial role in the REST algorithm (Sec. 3.2), by enabling the algorithm to simultaneously consider an entire family of term orderings during the exploration of rewrite paths. In this section, we provide a formal definition for ordering constraint algebras and describe the construction of an algebra for the RPQO.

► **Definition 14** (Ordering Constraint Algebra). *An Ordering Constraint Algebra (OCA) $\mathcal{A}_{(T,\Gamma)}$ over a set of terms T and term ordering family Γ , is a five-tuple $\mathcal{A}_{(T,\Gamma)} \doteq \langle C, \gamma, \top, \text{refine}, \text{sat} \rangle$, where:*

1. C , the constraint language, can be any non-empty set. Elements of C are called constraints, and are ranged over by c .
2. γ , the concretization function of $\mathcal{A}_{(T,\Gamma)}$, is a function from elements of C to subsets of Γ .
3. \top , the top constraint, is a distinguished constant from C , satisfying $\gamma(\top) = \Gamma$.
4. refine , the refinement function, is a function $C \rightarrow T \rightarrow T \rightarrow C$, satisfying (for all c, t_l, t_r) $\gamma(\text{refine}(c, t_l, t_r)) = \{ \succ \mid \succ \in \gamma(c) \wedge t_l \succ t_r \}$.
5. sat , the satisfiability function, is a function $C \rightarrow \text{Bool}$, satisfying (for all c) $\text{sat}(c) = \text{true} \Leftrightarrow \gamma(c) \neq \emptyset$.

The functions \top , refine , and sat are all called from our REST algorithm (Figure 3), and must be implemented as (terminating) functions when implementing REST. Specifically, REST instantiates the initial path with constraints $c = \top$. When a path can be extended via a rewrite application $t_l \rightarrow_R t_r$, REST refines the prior path constraints c to $c' \doteq \text{refine}(c, t_l, t_r)$. Then, the new term is added to the path only if the new constraints are satisfiable ($\text{sat}(c')$ holds); that is, if c' admits an ordering that orients the generated path. The function γ need *not* be implemented in practice; it is purely a mathematical concept used to give semantics to the algebra.

Given terms T and a finite term ordering family Γ , a trivial OCA is obtained by letting $C = \mathcal{P}(\Gamma)$, and making γ the identity function; straightforward corresponding elements \top , refine , and sat can be directly read off from the constraints in the definition above.

However, for efficiency reasons (or in order to support potentially infinite sets of orderings, which our theory allows), tracking these sets symbolically via some suitably chosen constraint language can be preferable. For example, consider lexicographic orderings on pairs of constants, represented by a set T of terms of the form $p(q_1, q_2)$ for a fixed function symbol p and q_1, q_2 chosen from some finite set of constant symbols Q . We choose the term ordering family $\Gamma = \{ \succ_{\text{lex}(\succ)} \mid \succ \text{ is a total order on } Q \}$ writing $\succ_{\text{lex}(\succ)}$ to mean the corresponding lexicographic ordering on $p(q_1, q_2)$ terms generated from an ordering \succ on Q .

A possible OCA over these T and Γ can be defined by choosing the constraint language C to be *formulas*: conjunctions and disjunctions of atomic constraints of the forms $q_1 > q_2$ and $q_1 = q_2$ prescribing conditions on the underlying orderings on Q . The concretization γ is given by $\gamma(c) = \{ \succ_{\text{lex}(\succ)} \mid \succ \text{ satisfies } c \}$, i.e., a constraint maps to all lexicographic orders generated from orderings of Q that satisfy the constraints described by c , defined in the natural way. We define \top to be e.g., $q = q$ for some $q \in Q$. A satisfiability function sat can be implemented by checking the satisfiability of c as a formula. Finally, by inverting the standard definition of lexicographic ordering, we define:

$$\text{refine}(c, p(q_1, q_2), p(r_1, r_2)) = c \wedge (q_1 > r_1 \vee (q_1 = r_1 \wedge q_2 > r_2))$$

489 Using this example algebra, suppose that REST explores two potential rewrite steps
 490 $p(a_1, a_2) \rightarrow p(b_1, a_2) \rightarrow p(a_1, a_1)$. Starting from the initial constraint $c_0 = \top$, the con-
 491 straint for the first step $c_1 \doteq \text{refine}(c_0, p(a_1, a_2), p(b_1, a_2)) = a_1 > b_1 \vee (a_1 = b_1 \wedge a_2 > a_2)$ is
 492 satisfiable, e.g., for any total order for which $a_1 > b_1$. However, considering the subsequent
 493 step, the refined constraint $c_2 \doteq \text{refine}(c_1, p(b_1, a_2), p(a_1, a_1))$, computed as $c_2 = c_1 \wedge (a_2 >$
 494 $a_2 \vee (a_2 = a_2 \wedge b_1 > a_1))$ is no longer satisfiable. Note that this allows us to conclude that
 495 there is no lexicographic ordering allowing this sequence of two steps, even without explicitly
 496 constructing any orderings.

497 We now describe an OCA for RPQOs (Sec. 4.1.2), based on a compact representation of
 498 sets of these orderings.

499 4.2.1 An Ordering Constraint Algebra for \succ_{rpo}

500 The OCA for RPQOs enables their usage in REST's proof search. One simple but computa-
 501 tionally intractable approach would be to enumerate the entire set of RPQOs that orient a
 502 path; continuing the path so long as the set is not empty. This has two drawbacks. First,
 503 the number of RPQOs grows at an extremely fast rate with respect to the number of func-
 504 tion symbols; for example there are 6,942 RPQOs describing five function symbols, and
 505 209,527 over six. Second, most of these orderings differ in ways that are not relevant to the
 506 comparisons made by REST.

507 Instead, we define a language to succinctly describe the set of candidate RPQOs, by
 508 calculating the minimal constraints that would ensure orientation of the path of terms;
 509 REST continues so long as there is some RPQO that satisfies the constraints. Crucially the
 510 satisfiability check can be performed effectively using an SMT solver, as described in Sec. 6.2,
 511 without actually instantiating any orderings.

512 Before formally describing the language, we begin with some examples, showing how the
 513 ordering constraints could be constructed to guide the termination check of REST.

► **Example 15** (Satisfiability of Ordering Constraints). Consider the following rewrite path
 given by the rules $r_1 \doteq f(g(x), y) \rightarrow g(f(y, y))$ and $r_2 \doteq f(x, x) \rightarrow f(k, x)$:

$$f(g(h), k) \rightarrow_{r_1} g(f(h, h)) \rightarrow_{r_2} g(f(k, h))$$

514 To perform the first rewrite REST has to ensure that there exists an RPQO \succ_{rpo} such
 515 that $f(g(h), k) \succ_{rpo} g(f(h, h))$. Following from Definition 10, we obtain three possibilities:

- 516 1. $f >_{\mathcal{F}} g$ and $\{f(g(h), k)\} >_{M(rpo)} \{f(h, h)\}$, or
- 517 2. $g >_{\mathcal{F}} f$ and $\{g(h), k\} \succ_{M(rpo)} \{g(f(h, h))\}$, or
- 518 3. $f \approx g$ and $\{g(h), k\} \succ_{M(rpo)} \{f(h, h)\}$.

We can further simplify these using the definition of the multiset quasi-ordering (Def. 9).
 Concretely, the multiset comparison of (1) always holds, while the multiset comparisons of
 (2) and (3) reduce to $k >_{\mathcal{F}} f \wedge k >_{\mathcal{F}} g \wedge k >_{\mathcal{F}} h$. Thus, we can define the exact constraints
 c_0 on \succ_{rpo} to satisfy $f(g(h), k) \succ_{rpo} g(f(h, h))$ as

$$c_0 \doteq f >_{\mathcal{F}} g \vee (k >_{\mathcal{F}} f \wedge k >_{\mathcal{F}} g \wedge k >_{\mathcal{F}} h)$$

519 Since there exist many quasi-orderings satisfying this formula (trivially, the one containing
 520 the single relation $f >_{\mathcal{F}} g$), the first rewrite is satisfiable.

521 Similarly, for the second rewrite, the comparison $g(f(z, z)) \succ_{rpo} g(f(k, z))$ entails the
 522 constraints $c_1 \doteq z \succ_{\mathcal{F}} k$. To perform this second rewrite the conjunction of c_0 and c_1
 523 must be satisfiable. Since the second disjunct of c_0 contradicts c_1 , the resulting constraints
 524 $f >_{\mathcal{F}} g \wedge z \succ_{\mathcal{F}} k$ is satisfiable by an RPQO, thus the path is satisfiable.

► **Example 16** (Unsatisfiable Ordering Constraint). As a second example, consider the rewrite rules $r_1 \doteq f(x) \rightarrow g(s(x))$ and $r_2 \doteq g(s(x)) \rightarrow f(h(x))$. These rewrite rules can clearly cause divergence, as applying rule r_1 followed by r_2 will enable a subsequent application of r_1 to a larger term. Now let's examine how our ordering constraint algebra can show the unsatisfiability of the diverging path:

$$f(z) \rightarrow_{r_1} g(s(z)) \not\rightarrow_{r_2} f(h(z))$$

525 $f(z) \succ_{rpo} g(s(z))$ requires $c_0 \doteq f > g \wedge f > s$ which is satisfiable, but $g(s(z)) \succ_{rpo} f(h(z))$
 526 requires $c_1 \doteq (g \geq f \wedge g \geq h) \vee (g \geq f \wedge s \geq h) \vee (s > f \wedge s > h)$, which, although satisfiable
 527 on it's own, conflicts with c_0 . Since no *RPQO* can satisfy both c_0 and c_1 , the rewrite path
 528 is not satisfiable.

Having primed intuition through the examples, we now present a way to compute such constraints. First, it is clear that we can define an RPQO based on the precedence over symbols \mathcal{F} . Therefore, we define our language of constraints to include the standard logical operators as well as atoms representing the relations between elements of \mathcal{F} , as:

$$C_{\mathcal{F}} \doteq f >_{\mathcal{F}} g \mid f \approx g \mid C_{\mathcal{F}} \wedge C_{\mathcal{F}} \mid C_{\mathcal{F}} \vee C_{\mathcal{F}} \mid \top \mid \perp$$

529 Next, we lift our definition of *RPQO* and the multiset quasi-ordering to derive functions:
 530 $rpo : \mathcal{T} \rightarrow \mathcal{T} \rightarrow C_{\mathcal{F}}$, and $mul : (\mathcal{T} \rightarrow \mathcal{T} \rightarrow C_{\mathcal{F}}) \rightarrow M(\mathcal{T}) \rightarrow M(\mathcal{T}) \rightarrow C_{\mathcal{F}}$. rpo is derived by
 531 a straightforward translation of Def. 10:

$$532 \quad rpo(f(t_1, \dots, t_m), g(u_1, \dots, u_n)) = \begin{aligned} & f >_{\mathcal{F}} g \quad \wedge \quad mul'(rpo, \{f(t_1, \dots, t_m)\}, \{u_1, \dots, u_n\}) \vee \\ & g >_{\mathcal{F}} f \quad \wedge \quad mul(rpo, \{t_1, \dots, t_m\}, \{g(u_1, \dots, u_n)\}) \vee \\ & f \approx g \quad \wedge \quad mul(rpo, \{t_1, \dots, t_m\}, \{u_1, \dots, u_n\}) \end{aligned}$$

533 where mul' is the strict multiset comparison: $mul'(f, T, U) = mul(f, T, U) \wedge \neg mul(f, U, T)$.
 534 $\neg : C_{\mathcal{F}} \rightarrow C_{\mathcal{F}}$ inverts the constraints, with $\neg(f >_{\mathcal{F}} g) = f \approx g \vee g >_{\mathcal{F}} f$ and $\neg(f \approx g) =$
 535 $f >_{\mathcal{F}} g \vee g >_{\mathcal{F}} f$; the other cases are defined in the typical way.

536 The definition for mul is more complex. Recall that $T \succ_{M(X)} U$ when U can be obtained
 537 from T by replacing zero or more elements in T with the same number of equal (with respect
 538 to \succ_X) elements, and by replacing zero or more elements in T with a finite number of smaller
 539 ones. Therefore each justification for $\{t_1, \dots, t_m\} \succ_{M(X)} \{u_1, \dots, u_n\}$ can be represented
 540 by a bipartite graph with nodes labeled t_1, \dots, t_m and u_1, \dots, u_n , such that:

- 541 1. Each node u_i has exactly one incoming edge from some node t_j .
- 542 2. If a node t_i has exactly one outgoing edge, it is labeled either **GT** or **EQ**.
- 543 3. If a node t_i has more than one outgoing edge, it is labeled **GT**.

544 $mul(f, \{t_1, \dots, t_m\}, \{u_1, \dots, u_n\})$ generates all such graphs: for each graph converts each
 545 labeled edge (t, u, \mathbf{EQ}) to the formula $f(t, u) \wedge f(u, t)$, each edge (t, u, \mathbf{GT}) to the formula
 546 $f(t, u) \wedge \neg f(u, t)$, and finally joins the formulas for the graph via a conjunction. The resulting
 547 constraint is defined to be the disjunction of the formulas generated from all such graphs.

548 Having defined the lifting of recursive path quasi-orderings to the language of constraints,
 549 we define our ordering constraint algebra $\mathcal{A}_{(\mathcal{T}, \Gamma)}$ as the tuple $\langle C_{\mathcal{F}}, \top, refine, \gamma, sat \rangle$ where:

- 550 ■ $refine(c, t, u) = c \wedge rpo(t, u)$,
- 551 ■ Γ is the set of all RPQOs,
- 552 ■ $\gamma(c)$ is the set of RPQOs derived from the underlying quasi-orders $\succ_{\mathcal{F}}$ that satisfy c , and
- 553 ■ $sat(c) = true$ if and only if there exists a quasi-order $\succ_{\mathcal{F}}$ satisfying c .

554 That $\mathcal{A}_{(\mathcal{T}, \Gamma)}$ is an OCA, i.e., satisfies the requirements of Def. 14, follows by construction.
 555 Namely, the function $rpo(t, u)$ produces constraints c such that, for any RPQO \succ_{rpo} , $t \succ_{rpo} u$
 556 if and only if its underlying ordering $\succ_{\mathcal{F}}$ satisfies c . In Sec. 6.2 we further discuss how the
 557 satisfiability check is mechanized and implemented using an SMT solver.

558 Having shown that using RPQOs as a term ordering is useful for theorem proving, sat-
 559 isfies the necessary properties for REST, and admits an efficient ordering constraint algebra,
 560 we continue our formal work by stating and proving the metaproperties of REST.

561 **5 REST Metaproperties: Soundness, Completeness, and Termination**

562 We now present the metaproperties of the REST algorithm defined in Figure 3. We show
 563 correctness (Theorem 17), completeness (Theorem 19) relative to the input term ordering
 564 family (recall that its instances must all be thin well-quasi-orderings), and termination
 565 (Theorem 22) which requires that calls to the OCA functions used in the algorithm, as well
 566 as the external oracle function, themselves terminate. The property that the orderings are
 567 thin well-founded guarantees in particular that any *duplicate-free* path (such as those that
 568 REST generates) that can be oriented by any of these orderings is guaranteed to be finite.
 569 We provide here the key invariants and statements of the formal results, and relegate the
 570 detailed proofs to App. A.

571 **5.1 Formal Definitions**

572 Our formalism of rewriting is standard; based on the terminology of [28]. Our language
 573 consists of the following:

- 574 1. An infinite set of meta-variables (the variables for rewrite rules) \mathcal{V} with elements $X, Y,$
 575 \dots
- 576 2. A finite set of function symbols \mathcal{F} with elements f, g, \dots, x, y, \dots
 577 Each operator is associated with a fixed numeric arity and types for its arguments and
 578 result (elided here, for simplicity). By convention, we use the variables x, y to range over
 579 zero-arity function symbols (constants).
- 580 3. A set of terms \mathcal{T} with elements t, u, \dots inductively defined as follows: (a) $X \in \mathcal{V} \Rightarrow X \in \mathcal{T}$
 581 and (b) $f \in \mathcal{F}$, f has arity n , $t_1, \dots, t_n \in \mathcal{T} \Rightarrow f(t_1, \dots, t_n) \in \mathcal{T}$.

582 We use $FV(t)$ to refer to the set of meta-variables in t . A term t is *ground* if $FV(t) = \emptyset$.

583 A *substitution* $\sigma \subseteq \mathcal{V} \times \mathcal{T}$ is a mapping from meta-variables to terms. We write $\sigma \cdot t$ to
 584 denote the simultaneous application of the substitution: namely, $\sigma \cdot t$ replaces each occurrence
 585 of each meta-variable X in t with $\sigma(X)$. A substitution σ *grounds* t if, for all $X \in FV(t)$,
 586 $\sigma(X)$ is a ground term. A substitution σ *unifies* two terms t and u if $\sigma \cdot t = \sigma \cdot u$.

587 A *context* E is a term-like object that contains exactly one term placeholder \bullet . If t is a
 588 term, then $E[t]$ is the term generated by replacing the \bullet in E with t .

589 A *rewrite rule* r is a pair of terms $r \doteq (t, u)$ such that $FV(u) \subseteq FV(t)$ and $t \notin \mathcal{V}$. Each
 590 rewrite rule $r \doteq (t, u)$ defines a binary relation \rightarrow_r which is the smallest relation such that,
 591 for all contexts E and substitutions σ grounding t (and therefore u), $E[\sigma \cdot t] \rightarrow_r E[\sigma \cdot u]$.

592 We use R to range over sets of rewrite rules. We write $v \rightarrow_R w$ iff $v \rightarrow_r w$ for some
 593 $r \in R$.

594 For oracle functions (from terms to sets of terms) \mathcal{E} , we write $t \rightarrow_{\mathcal{E}} t'$ iff $t' \in \mathcal{E}(t)$. We
 595 write $t \rightarrow_{R+\mathcal{E}} t'$ if $t \rightarrow_R t'$ or $t \rightarrow_{\mathcal{E}} t'$. For a relation \rightarrow we write \rightarrow^* for its reflexive,
 596 transitive closure. A *path* is a list of terms. A binary relation \succ *orients* a path t_1, \dots, t_n if
 597 $\forall i, 1 \leq i < n, t_i \succ t_{i+1}$.

598 **5.2 Soundness**

599 Soundness of REST means that any term of the output ($u \in \text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$) can be derived
600 from the original input term by some combination of term rewriting steps from R and steps
601 via the oracle function \mathcal{E} (in other words, $t_0 \rightarrow_{R+\mathcal{E}}^* u$).

602 Our proof relies on the following simple invariant of REST: any path stored in the stack
603 during the execution of the algorithm can be derived by the rewrite rules in R or the external
604 oracle \mathcal{E} .

605 ► **REST Invariant 1** (Path Invariant). *For any execution of $\text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$, at the start of*
606 *any iteration of the main loop, for each $(ts, c) \in p$, the list ts is a path of $R + \mathcal{E}$ starting*
607 *from t_0 .*

608 **Proof.** (Sketch:) By straightforward induction on iterations of the main loop. ◀

609 ► **Theorem 17** (Soundness of REST). *For all R , u , and t_0 , if $u \in \text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$, then*
610 *$t_0 \rightarrow_{R+\mathcal{E}}^* u$.*

611 **Proof.** In each iteration of REST, the term t added to the output o is the last element of the
612 list ts for the tuple $(ts, c) \in p$. By Invariant 1, t must be on the path of $R + \mathcal{E}$ starting from
613 t_0 . ◀

614 **5.3 Completeness**

615 A naïve completeness statement for REST might be that, for any terms t_0 and u , if $t_0 \rightarrow_{R+\mathcal{E}}^* u$
616 then u is in our output ($u \in \text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$). This result doesn't hold in general by design,
617 since REST explores only paths permitted by at least one candidate instance of its input term
618 ordering family. We prove this *relative* completeness result in two stages. First (Theorem 18),
619 we show that completeness always holds if all steps only involve the external oracle. Then
620 (Theorem 19), we prove relative completeness of REST with respect to the provided term
621 ordering family. We begin by stating another simple invariant of our algorithm: that any
622 term appearing in a path in the stack p , will belong to the final output:

623 ► **REST Invariant 2.** *For any execution of $\text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$, at the start of any iteration of*
624 *the main loop, if $t \in ts$ and $(ts, c) \in p$, then, when the algorithm terminates, we will have*
625 *$t \in \text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$.*

626 **Proof.** (Sketch:) We can prove inductively that terms contained in any list in p either
627 remain in p or end up in o ; since p is empty on termination, the result follows. ◀

628 ► **Theorem 18** (Completeness w.r.t. \mathcal{E}). *For all R , u , and t_0 , if $t_0 \rightarrow_{\mathcal{E}}^* u$, then $u \in$*
629 *$\text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$.*

630 **Proof.** (Sketch:) Since, \mathcal{E} is strongly normalizing, the path of terms $t_0 \rightarrow_{\mathcal{E}} \dots \rightarrow_{\mathcal{E}} u$ will
631 not contain any duplicates; REST will therefore insert each term in the path into ts . Since u
632 is in that path, Invariant 2 ensures $u \in \text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$. ◀

633 ► **Theorem 19** (Relative Completeness). *For all R , u , and t_0 , if $t_0 \rightarrow_{R+\mathcal{E}}^* u$ and there exists*
634 *an ordering $\succ \in \gamma(\mathbb{T})$ that orients the path justifying $t_0 \rightarrow_{R+\mathcal{E}}^* u$, then $u \in \text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$.*

635 **Proof.** (Sketch:) The proof structure is similar to Theorem 18; in this case the terms in the
636 path are guaranteed to be in ts because some ordering in $\gamma(\mathbb{T})$ can orient the path. ◀

637 **5.4 Termination**

638 Termination of REST requires appropriate conditions on the external oracle \mathcal{E} and the order-
 639 ing constraint algebra \mathcal{A} employed. We formally define these requirements and then prove
 640 termination of REST.

641 **► Definition 20** (Well-Founded ordering constraint algebras). *For ordering constraint algebras*
 642 $\mathcal{A} = \langle C, \top, \text{refine}, \text{sat}, \gamma \rangle$, for $c, c' \in C$, we say c' strictly refines c (denoted $c' \sqsubset_{\mathcal{A}} c$) if
 643 $c' = \text{refine}(c, t, u)$ for some terms t and u , and $\gamma(c') \subset \gamma(c)$. Then, we say \mathcal{A} is well-founded
 644 if $\sqsubset_{\mathcal{A}}$ is.

645 Down every path explored by REST, the tracked constraint is only ever refined; well-foundedness
 646 of \mathcal{A} guarantees that finitely many such refinements can be strict.

647 We note that if the OCA describes a finite set of orderings, then it is trivially well-
 648 founded: \subset is well-founded on finite sets. For example, the ordering constraint algebra for
 649 RPQOs (Sec. 4.2.1) is well-founded when the set of function symbols \mathcal{F} is finite, as there
 650 are only a finite number of possible RPQOs over a finite set of function symbols.

651 **► Definition 21.** *A relation $t_l \rightarrow t_r$ is normalizing if it does not admit an infinite path and*
 652 *bounded if for each t_l it only admits finite t_r .*

653 **► Theorem 22** (Termination of REST). *For any finite set of rewriting rules R , if:*

- 654 1. $\rightarrow_{\mathcal{E}}$ is normalizing and bounded,
- 655 2. The refine and sat functions from \mathcal{A} are decidable (always-terminating, in an implemen-
 656 tation),
- 657 3. \mathcal{A} is well-founded,

658 then, for all terms t_0 , $\text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$ terminates.

659 **Proof.** (Sketch:) The paths constructed by REST implicitly constructs a finitely branching
 660 tree, and the four restrictions ensures that all paths down the tree are finite. This ensures
 661 that the resulting tree is finite; and thus that REST’s implicit construction of the tree will
 662 terminate. ◀

663 Note that any deterministic, terminating external oracle function satisfies the first re-
 664 quirement. Having completed the formalization, we now move on to the details of our
 665 implementation.

666 **6 Implementation of REST**

667 We implemented REST as a standalone library, comprising 2337 lines of Haskell code (Sec. 6.1).
 668 Our implementation includes the REST algorithm, several ordering constraint algebra imple-
 669 mentations (including RPQOs [Def. 10]) and exposes the API for implementing ordering
 670 constraint algebras (Sec. 6.2). We integrated this library into the Liquid Haskell program
 671 verifier [49] (Sec. 6.3), where we chose the task of applying *lemmas* in Liquid Haskell proofs
 672 as a suitable target problem for automation via REST.

```

-- Interface of OC Algebra
data OC C T = OC
  { top    :: C
  , refine :: C → T → T → C
  , sat    :: C → IO Bool
  }

-- Language of Logical Formulas
data LF A = LTrue | LFalse
  | A >: A | A :=: A
  | LF A :^: LF A | LF A :V: LF A

-- Implementation of OC Algebra
rpoOC :: OC (LF F) T
rpoOC = OC LTrue refine sat where

  refine :: LF F → T → T → LF F
  refine c t u =
    c :^: rpo t u -- As in Def 10

  sat :: LF F → IO Bool
  sat = smtSat . toSMT -- SMT Interface

```

■ **Figure 5** The implementation of our RPQO Ordering Constraint Algebra

6.1 The REST Library

Our REST implementation is developed in Haskell and can be used directly by other Haskell projects. The library is designed modularly; for example, a client of the library can decide to use REST only for comparing terms via an OCA, without also using the proof search algorithm of Sec. 3.2. In addition, our library has a small code footprint and can be used with or without external solvers, making it ideal for integration into existing program analysis tools and theorem provers.

Furthermore, we include in the library built-in helper utilities for encoding and solving constraints on term orderings. Although the library enables integration of arbitrary solvers; it provides several built-in solvers for constraints on finite WQOs and also provides an interface for solving constraints with external SMT solvers. These utilities comprise the majority of the code in the REST library (1369 out of the 2337 lines).

Our implementation defines the OCA interface of Sec. 4.2 and provides three built-in instances for RPQOs, LPQOs (derived from the Lexicographic path ordering), and KBQOs (Sec. 4.1.1). The helper utilities included in the library enable a concise implementation of these OCAs: the three OCA implementations consist of 200 lines of code in total.

To facilitate debugging and evaluation of OCAs, the library also provides a standalone executable that produces visualizations of the rewrite paths that REST explores when using the OCA to compute the rewrites paths from a given term. Figure 4 and Figure 8 were produced using this functionality; we also note that the visualization is also capable of displaying the accumulated constraints on the ordering at each node in the tree.

We now describe the interface for defining OCAs in our REST implementation, via a presentation of the RPQO algebra in the library.

6.2 Efficient Implementations of OCAs in REST

Figure 5 presents REST’s library interface for ordering constraint algebras and the implementation of RPQOs. The interface `oc` is parametric in the language of constraints `c` and the type of terms `t`. The logical formulas `LF A` describe constraints on WQOs over `A`, in the case of RPQOs, `LF F` tracks constraints on the underlying precedence of function symbols.

Our implementation `rpoOC` defines the initial constraints `top` to be `LTrue`, (intuitively, permitting any RPQO). The function `refine c t u` conjoins the current constraints `c` with the constraints `rpo t u`, ensuring $t \succcurlyeq u$. Finally the `sat` function converts the constraints into an equisatisfiable SMT formula, by encoding each distinct function symbol as an SMT integer variable, encoding the logical operators as their SMT equivalent, and checking for

```

{-@ example1 :: s0 : Set → { s1 : Set | IsDisjoint s0 s1 } → f : (Set → a) → { f ((s0
  \\/ s1) /\ s0) = f s0 } @-}
example1 :: Set → Set → (Set → a) → Unit
example1 s0 s1 f =
  f ((s0 \\/ s1) /\ s0)          ? distribUnion s0 s1 s0
=== f ((s0 /\ s0) \\/ (s1 /\ s0)) ? idemInter s0
=== f (s0 \\/ (s1 /\ s0))        ? symmInter s1 s0
=== f (s0 \\/ (s0 /\ s1))        -- Disjoint
=== f (s0 \\/ emptySet)         ? emptyUnion s0
=== f s0
*** QED

```

■ **Figure 6** Liquid Haskell version of the proof from Example 1.

706 satisfiability of the resulting formula.

707 REST’s interface supports arbitrary implementations for ordering constraints and is not
 708 dependent on any particular ordering, constraint language, or solver. For example, the `sat`
 709 function for RPQOs could evaluate the formulas using an alternative solver; in fact REST
 710 includes a built-in solver for this purpose, although it does not achieve as high performance
 711 as the SMT-based approach.

712 6.3 Integration of REST in Liquid Haskell

713 We used REST to automate lemma application in Liquid Haskell. Here we provide a brief
 714 overview of Liquid Haskell (Sec. 6.3.1), how REST is used to automate lemma instantiations
 715 (Sec. 6.3.2) and how it mutually interacts with the existing Liquid Haskell automation
 716 (Sec. 6.3.3).

717 6.3.1 Liquid Haskell and Program Lemmas

718 Liquid Haskell performs program verification via *refinement types* for Haskell; function types
 719 can be annotated with refinements that capture logical/value constraints about the func-
 720 tion’s parameters, return value and their relation. For example, the function `example1`
 721 in Figure 6 ports the set example of Example 1 to Liquid Haskell, without any use of REST.
 722 User-defined lemmas amount to nothing more than additional program functions, whose
 723 refinement types express the logical requirements of the lemma. The first line of the figure
 724 is special comment syntax used in Liquid Haskell to introduce refinement types; it expresses
 725 that the first parameter `s0` is unconstrained, while the second `s1` is refined in terms of `s0`: it
 726 must be some value such that `IsDisjoint s0 s1` holds. The refinement type on the (unit)
 727 return value expresses the proof goal; the body of the function provides the proof of this
 728 lemma. The proof is written in equational style; the `?` annotations specify lemmas used to
 729 justify proof steps [48]. The penultimate step requires no lemma; the verifier can discharge
 730 it based on the refinement on the `s1` parameter.

731 Lemmas already proven can be used in the proof of further lemmas; as is standard for
 732 program verification, care needs to be taken to avoid circular reasoning. Liquid Haskell
 733 ensures this via well-founded recursion: lemmas can only be instantiated recursively with
 734 smaller arguments.

6.3.2 REST for Automatic Lemma Application in Liquid Haskell

We apply REST to automate the application of equality lemmas in the context of Liquid Haskell. The basic idea is to extract a set of rewrite rules from a set of refinement-typed functions, each of which must have a refinement type signature of the following shape:

```
{-@ rrule :: x1:t1 → ... → xn:tn → {v:() | e_l = e_r } @-}
```

In particular, the equality $e_l = e_r$ refinement of the (unit) return value generates potential rewrite rules to feed to REST, in both directions. Let $FV(e)$ be the free variables of e , if $FV(e_r) \subseteq FV(e_l)$ and $e_l \notin \{x_1, \dots, x_n\}$ then $e_l \rightarrow e_r$ is generated as a rewrite rule. Symmetrically, if $FV(e_l) \subseteq FV(e_r)$ and $e_r \notin \{x_1, \dots, x_n\}$ then $e_r \rightarrow e_l$ is generated as a rewrite rule. These rewrite rules are fed to REST along with the current terms we are trying to equate in the proof goal; any rewrites performed by REST are fed back to the context of the verifier as assumed equalities.

Since the extracted rewrite rules are defined as refinement-typed expressions, our implementation technically goes beyond simple term rewriting, since instantiations of these rules in our implementation are also refinement-type-checked; i.e., it instantiates only the rules with expressions of the proper refined type, achieving a form of conditional rewriting [27].

Selective Activation of Lemmas: Local and Global Rewrite Rules In our Liquid Haskell extension, the user can activate a rewrite rule globally or locally, using the `rewrite` and `rewriteWith` pragmas, *resp.* For example, with the below annotations

```
{-@ rewrite global @-}
{-@ rewriteWith theorem [local] @-}
```

the rule `global` will be active when verifying every function in the current Haskell module, while the rule `local` is used only when verifying `theorem`.


Preventing Circular Reasoning Our implementation finally ensures that rewrites cannot be used to justify circular reasoning, by checking that there are no cycles induced by our `rewrite` and `rewriteWith` pragmas. For example, the below, unsound, circular dependency will be rejected with a rewrite error by our implementation.

```
{-@ rewriteWith p1 [p2] @-}
{-@ rewriteWith p2 [p1] @-}
{-@ p1, p2 :: x:Int → { x = x + 1 } @-}
p1 _ = () ; p2 = p1
```

To prevent circular dependencies, we check that the dependency graph of the rewrite rules (which are made available for proving with) has no cycles. This simple restriction is stronger than strictly necessary; a more-complex termination check could allow rewrites to be mutually justified by ensuring that recursive rewrites are applied with smaller arguments. In practice, our coarse check isn't too restrictive: because Haskell's module system enforces acyclicity of imports, rewrite rules placed in their own module can be freely referenced by importing the library.

Lemma Automation Using our implementation, the same Example 1 proven manually in Figure 6 can be alternatively proven (with all relevant rewrite rules in scope) as follows:

```
{-@ example1 :: s0 : Set → { s1 : Set | IsDisjoint s0 s1 } → f : (Set → a) →
    { f ((s0 \ / s1) /\ s0) = f s0 } @-}
example1 s0 s1 _ = ()
```



■ **Figure 7** Interaction between PLE and REST.

786 The proof is fully automatic: no manual lemma calls are needed as these are all handled
 787 by REST. Integrating REST into Liquid Haskell required around 500 lines of code, mainly for
 788 surface syntax.

789 6.3.3 Mutual PLE and REST Interaction

790 Liquid Haskell includes a technique called *Proof by Logical Evaluation* (PLE) [50] for au-
 791 tomating the expansion of terminating program function definitions. PLE expands function
 792 calls into single cases of their (possibly conditional) bodies exactly when the verifier can
 793 prove that a unique case definitely applies. This check is performed via SMT and so can
 794 condition on arbitrary logical information; in our implementation, this forms a natural com-
 795 plement to the term rewriting of REST, and plays the role of its external oracle (cf. Sec. 3).
 796 Since PLE is proven terminating [50], the termination of this collaboration is also guaranteed
 797 (cf. Sec. 5).

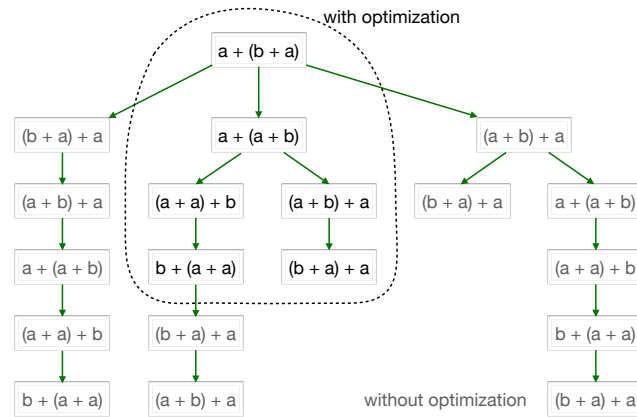
798 Figure 7 summarizes the mutual interaction between PLE and REST on a verification
 799 condition $\Phi \vdash p$, where Φ is an environment of assumptions. PLE also takes as input a
 800 set \mathcal{F} of (provably) terminating, user-defined function definitions that it iteratively evalu-
 801 ates. Meanwhile, REST is provided with the rewrite rules extracted from in-scope lemmas in
 802 the program (cf. Sec. 6.3.2); these two techniques can then generate paths of equal terms in-
 803 cluding steps justified by each technique. For example, consider the following simple lemma
 804 `countPosExtra`, stating that the number of strictly positive values in `xs ++ [y]` is the number
 805 in `xs`, provided that `y <= 0`, and a lemma stating that `countPos` of two lists appended gives
 806 the same result if their orders are swapped.

```

807 {-@ lm :: xs : [Int] -> ys : [Int] -> { countPos (xs ++ ys) = countPos (ys ++ xs) } @-}
808
809
810 {-@ rewriteWith countPosExtra [lm] @-}
811 {-@ countPosExtra :: xs : [Int] -> { y : Int | y <= 0 } ->
812     { countPos (xs ++ [y]) = countPos xs } @-}
813 countPosExtra :: [Int] -> Int -> ()
814 countPosExtra _ _ = () -- proof is fully automatic!
815

```

816 The proof requires rewriting `countPos(xs ++ [y])` first via lemma `lm` (by REST), expanding
 817 the definition of `++` twice (via PLE) to give `countPos(y:xs)`, and finally one more PLE step
 818 evaluating `countPos`, using the logical fact that `y` is not positive. Note in particular that the
 819 first step requires applying an external lemma (out of scope for PLE) and the last requires
 820 SMT reasoning not expressible by term rewriting. The two techniques together allow for a
 821 fully automatic proof.



■ **Figure 8** Associative-commutative rewrites of $a + (b + a)$ generated by REST. Paths explored by REST with the explored terms optimization are within the dashed line. Using the explored terms optimization, REST only considers each term once.

6.4 Further Optimizing the REST Algorithm

When a rewrite system is branching, REST may encounter different rewrite paths from an initial term t to an arbitrary term u . For example, in Figure 8, the term $(b + a) + a$ is explored in 5 different paths. In general, REST cannot always ignore the repeat encounters of u , as a new path from t to u may impose ordering constraints enabling more rewrites in the future. Nonetheless, reducing the number of explored paths naturally improves performance. Therefore, we optimize REST based on the following observations:

1. A term t does not need to be revisited if all of its rewrites have already been visited.
2. If a term t was previously visited at constraints c , revisiting t at constraints c' is not necessary if c permits all orderings permitted by c' , i.e., $\gamma(c') \subseteq \gamma(c)$.

To implement this optimization, REST maintains a mapping M from terms to the logical constraints c each term was explored with (initially mapping all terms to **top**). To explore a term t under logical constraints c , the algorithm checks that this term is *explorable*, formally defined by:

$$\text{explorable}(t, c) \doteq t \notin M \vee (\neg(c \Rightarrow M[t]) \wedge \exists u. (t \rightarrow_R u \wedge \text{explorable}(u, c)))$$

This predicate ensures that either this term was not explored before or it comes with weaker constraints that can derive at least one new term in the path.

After exploring a new term, REST weakens the mapping M for this term to the disjunction of the constraints under which it was newly explored and those previously mapped to in M . With this optimization, a term will appear in more than one path in the REST graph only when it can lead to different terms in the path. This optimization critically reduces the number of explored terms even for small examples: as shown in Figure 8 where 19 vertices of the REST graph shown reduced to only the 6 in the dotted region.

7 Evaluation

Our evaluation seeks to answer three research questions:

§ 7.1: How does REST compare to existing rewriting tactics?

Property	LH+	Coq	Agda	Lean	Isabelle	Zeno	Isa+
Diverge	OK	loop	loop	fail	loop	OK	OK
Plus AC	OK	loop	loop	fail	fail	OK	OK
Congruence	OK	OK	OK	OK	OK	fail	OK

■ **Table 1** Comparison of REST with existing theorem provers. LH+ is Liquid Haskell with rewriting. The potential outcomes are **OK** when the property is proved; **loop** when no answer is returned after 300 sec; and **fail** when the property cannot be proven. Isa+ is Isabelle/HOL with Sledgehammer.

843 § 7.2: How does REST compare to E-matching based axiomatization?

844 § 7.3: Does REST simplify equational proofs?

845 We evaluate REST using the Liquid Haskell implementation described in Sec. 6. In Sec. 7.1,
 846 we compare our implementation’s rewriting functionality with that of other theorem provers,
 847 with respect to the challenges mentioned in Sec. 2. In Sec. 7.2, we compare against Dafny [33]
 848 by porting Dafny’s calculational proofs to Liquid Haskell, using rewriting to handle axiom
 849 instantiation. Finally, in Sec. 7.3, we port proofs from various sources into Liquid Haskell
 850 both with and without rewriting, and compare the performance and complexity of the
 851 resulting proofs.

852 7.1 Comparison with Other Theorem Provers

853 To compare REST with the rewriting functionality of other theorem provers, we developed
 854 three examples to test the five challenges described in Sec. 2 and compare our implemen-
 855 tation to that of other solvers. We chose to evaluate against Agda [39], Coq [11], Lean [5],
 856 Isabelle/HOL [38], and Zeno [44], as they are widely known theorem provers that either
 857 support a rewrite tactic, or use rewriting internally. Agda, Lean, and Isabelle/HOL allow
 858 user-defined rewrites. In Lean and Isabelle/HOL, the tactic for applying rewrite rules mul-
 859 tiple times is called `simp`; for simplification. Agda, Coq, and Isabelle/HOL’s implementation
 860 of rewriting can diverge for nonterminating rewrite systems [11, 1, 38]. On the other hand,
 861 Lean enforces termination, at least to some degree, by ensuring that associative and commu-
 862 tative operators can only be applied according to a well-founded ordering [4]. Zeno [44] does
 863 not allow for user-defined rewrite rules, rather it generates rewrites internally based on user-
 864 provided axioms. Sledgehammer [36, 42, 41] is a powerful tactic supported by Isabelle/HOL
 865 that (on top of the built-in rewriting) dispatches proof obligations to various external provers
 866 and succeeds when any of the external provers succeed; this tactic operates under a built-in
 867 (customizable) timeout.

868 1. **Diverge** tests how the prover handles the challenges 1 and 5: restricting the rewrite
 869 system to ensure termination and integrating external oracle steps. This example encodes
 870 a single (terminating) rewrite rule $f(x) \rightarrow g(s(s(x)))$ and terminating, mutually recursive
 871 function definitions for f and g . However, the combination of the rules and function expan-
 872 sions can cause divergence. This test also requires a simple proof that follows directly from
 873 the function definitions.

874 2. **Plus AC** tests the challenges 2 and 3 by encoding a task that requires a permissive term
 875 ordering. This example encodes p , q , and r , user-defined natural numbers, and requires that
 876 expressions such as $(p + q) + r$ can be rewritten into different groupings such as $(r + q) + p$,
 877 via associativity and commutativity rules.

878 3. **Congruence** is an additional test to ensure that the implementation of the rewrite
 879 system is permissive enough to generate the expected result. This test evaluates a basic

880 expected property, that the expressions $f(g(x))$ and $f(g'(x))$ can be proved equal if there
 881 exists a rewrite rule of the form $g(x) \rightarrow g'(x)$.

882 We present our results in Table 1. As expected, `Coq`, `Agda`, and `Isabelle/HOL` diverge on
 883 the first example, as they do not ensure termination of rewriting. `Lean` does not diverge, but
 884 it also fails to prove the theorem. Unsurprisingly, the commutativity axiom of `Plus AC` causes
 885 theorem provers that don't ensure termination of rewriting to loop. Although `Lean` ensures
 886 termination, it does not generate the necessary rewrite application in every case, because
 887 it orients associative-commutative rewriting applications according to a fixed order. With
 888 the exception of `Zeno`, all of the theorem provers tested were able to prove the necessary
 889 theorem for the final example. Our implementation succeeds on these three examples by
 890 implementing a permissive termination check based on non-strict orderings.

891 For this selection of simple but illustrative examples, the only tools to succeed on all
 892 cases are our implementation, and `Isabelle's Sledgehammer`. The latter combines a great
 893 many techniques which go beyond term rewriting. Nonetheless, we note that our novel
 894 approach provides a clear and general formal basis for incorporation with a wide variety of
 895 verifiers and reasoning techniques (due to its generic definition and formal requirements), and
 896 provides strong formal guarantees for such combinations. In particular, `REST` provides general
 897 termination and relative completeness guarantees, which `Sledgehammer` (via its timeout
 898 mechanism) does not.

899 7.2 Comparison with E-matching

900 To evaluate `REST` against the E-matching based approach to axiom instantiation, we com-
 901 pared with `Dafny` [33], a state-of-the-art program verifier. `Dafny` supports equational reason-
 902 ing via calculational proofs [34] and calculation with user-defined functions [2]. We ported
 903 the calculational proofs of [34] to `Liquid Haskell`, using rewriting to automatically instantiate
 904 the necessary axioms.

905 7.2.1 List Involution

906 Figure 9 shows an example taken directly from `Dafny` [34], proving that the reverse operation
 907 on lists is an involution, i.e., $\forall xs. reverse(reverse(xs)) = xs$. In this example, both `Liquid`
 908 `Haskell` and `Dafny` operate on inductively defined lists with user-defined functions `++` and
 909 `reverse`. The original `Dafny` proof goes through via the combination of a manual application
 910 of a lemma called `ReverseAppendDistrib` (stating that for all lists xs and ys , $reverse(xs ++$
 911 $ys) = reverse(ys) ++ reverse(xs)$) and induction on the size of the list.

912 Using term rewriting as enabled by our `REST` library, `Liquid Haskell` is able to simplify the
 913 proof, with `PLE` expanding the function definitions for `reverse` and `append`, and `REST` applying
 914 the necessary equality `reverse (reverse xs ++ [x]) = reverse [x] ++ reverse (reverse`
 915 `xs)`.

916 In `Dafny`, a similar simplification of the calculational proof is not possible; the proof
 917 fails if the manual equality steps are simply removed. We experimented further and found
 918 that the lemma `ReverseAppendDistrib` can be alternatively encoded as a user-defined axiom
 919 which, by itself, does not appear to cause trouble for E-matching, and with this change
 920 alone the proof succeeds without the need for this single lemma call. On the other hand, the
 921 equalities must still be mentioned for the calculational proof to succeed. Perhaps surprisingly,
 922 removing these intermediate equality steps caused `Dafny` to stall³; analysis with the `Axiom`

³ We include this version in App. D

```

lemma LemmaReverseTwice(xs: List)
  ensures reverse(reverse(xs)) == xs;
{
  match xs {
  case Nil =>
  case Cons(x, xrest) =>
    calc {
      reverse(reverse(xs));
      reverse(append(reverse(xrest), Cons(x, Nil)));
      { ReverseAppendDistrib(reverse(xrest), Cons(x, Nil)); }
      append(reverse(Cons(x, Nil)), reverse(reverse(xrest)));
      { LemmaReverseTwice(xrest); }
      append(reverse(Cons(x, Nil)), xrest);
      append(Cons(x, Nil), xrest);
      xs;
    }
  }
}

```

(a) Calculation-style proof in Dafny, from [34].

```

{-@ involutionP :: xs:[a] → {reverse (reverse xs) == xs } @-}
{-@ rewriteWith involutionP [distributivityP] @-}
involutionP [] = ()
involutionP (x:xs) = involutionP xs

```

(b) An equivalent proof implemented in Liquid Haskell extended with REST

■ **Figure 9** List Involution proofs in Liquid Haskell and Dafny

923 Profiler [7] indicated the presence of a (rather complex) matching loop involving the axiom
 924 `ReverseAppendDistrib` in combination with axioms internally generated by the verifier itself.
 925 This illustrates that achieving further automation of such E-matching-based proofs is not
 926 straightforward, and can easily lead to performance difficulties due to matching loops which
 927 can be hard to predict and understand, even in this state-of-the-art verifier. By contrast, REST
 928 can automatically provide the necessary equality steps for this proof without introducing
 929 any risk of non-termination.

930 7.2.2 Set Properties

931 Figure 10 shows the Dafny and Liquid Haskell proofs for the implication $s_0 \cap s_1 = \emptyset \implies$
 932 $f((s_0 \cup s_1) \cap s_0) = f(s_0)$.

933 Dafny uses a calculational proof to show the equality $(s_0 \cup s_1) \cap s_0 = s_0$, seemingly by
 934 applying distributivity. In fact, the distributivity aspect is not relevant to the proof; rather,
 935 the set equality in the proof syntax causes Dafny to instantiate the set extensionality axiom
 936 discharging the proof. It is for this reason that Dafny requires an extra proof step to prove
 937 $f((s_0 \cup s_1) \cap s_0) = f(s_0)$, as this term does not include an equality on sets, but rather on
 938 applications of f . Dafny's set axiomatization does not include the distributivity axiom, as
 939 such an axiom could easily lead to matching loops.

940 Using REST, it is safe to encode arbitrary lemmas as rewrite rules, as the termination is
 941 guaranteed; in this case the distributivity lemma can be used to complete the proof (and is
 942 permitted as a rewrite rule with the precedence $\cap > \cup$).

943 In conclusion, we have shown that using REST to apply rewrites could be used as an

```

lemma Proof<a>(s0: set<int>, s1: set<int>, f: set<int> → a)
  requires s0 * s1 == {}
  ensures f((s0 + s1) * s0) == f(s0) {
    calc {
      (s0 + s1) * s0; (s0 * s0) + (s1 * s0);
      s0;
    }
  }
}

```

(a) Proof in Dafny using built-in set axiomatization

```

{-@ assume unionEmpty :: ma : Set → {v : () | ma ∨ emptySet = ma } @-}
{-@ assume intersectComm :: ma : Set → mb : Set → {v : () | ma ∧ mb = mb ∧ ma } @-}
{-@ assume intersectSelf :: s0 : Set → { s0 ∧ s0 = s0 } @-}
{-@ assume unionIntersect :: s0 : Set → s1 : Set → s2 : Set →
      { (s0 ∨ s1) ∧ s2 = (s0 ∧ s2) ∨ (s1 ∧ s2) } @-}
{-@ rwDisjoint :: s0 : Set → {s1 : Set | IsDisjoint s0 s1} → { s0 ∧ s1 = emptySet } @-}

{-@ example1 :: s0 : Set → { s1 : Set | IsDisjoint s0 s1 } → f : (Set → a) →
      { f ((s0 ∨ s1) ∧ s0) = f s0 } @-}
example1 s0 s1 _ = ()

```

(b) An equivalent proof implemented in Liquid Haskell, with a user-defined axiomatization of sets.

■ **Figure 10** Set Proofs in Liquid Haskell and Dafny

944 alternative to E-matching based axiomatization. Furthermore, the termination guarantee
 945 of REST enables axioms that may give rise to matching loops to, instead, be encoded as
 946 rewrite rules.

947 7.3 Simplification of Equational Proofs

948 Finally, we evaluate how REST can simplify equational proofs. We chose to include the set
 949 example from [34] (described in Sec. 7.2.2), data structure proofs from [48], examples from
 950 the Liquid Haskell test suite, as well as our own case study. We developed each example
 951 in Liquid Haskell both with and without rewriting, and compared the timing and proof
 952 complexity. Each proof using rewriting was evaluated using each different ordering constraint
 953 algebras built-in to our Haskell REST library. The proofs in [48] were selected because they
 954 require induction, expansion of user-defined functions, and equational reasoning steps to
 955 prove properties about trees and lists. The examples from the Liquid Haskell test suite were
 956 taken to evaluate the rewriting across a range of representative proofs.

957 Our DSL case study evaluates the performance of our implementation using a larger set
 958 of rewrite rules, by verifying optimizations for a simple programming language, contain-
 959 ing statements (i.e., print, sequence, branches, repeats and no-ops) and expressions (i.e.,
 960 constants, variables, arithmetic and boolean expressions) using 23 rewrite rules. Our rewrit-
 961 ing technique can prove the kind of equivalences used in techniques such as supercompila-
 962 tion [8, 52, 46], by encoding the basic equality axioms as rewrite rules and using them
 963 to prove more complicated theorems. A full list of the axioms and proved theorems are
 964 available in App. C. We note that we encoded arithmetic operations as uninterpreted SMT
 965 functions, so that the built-in arithmetic theory of the SMT does not aid proof automation.

966 We present our results in Table 2. By using rewriting, we were able to eliminate all but
 967 two of the non-inductive axiom instantiations, while maintaining a reasonable verification

Name	Orig.	Cut	Rules	Time				
				Orig.	RPQO	LPQO	KBQO	Fuel
Set-Dafny	4	4	5	1.11s	✓1.15s	✓1.19s	✗1.13s	✓1.22s
Set-Mono	7	7	4	1.16s	✗1.40s	✗1.41s	✓1.47s	✓1.60s
List	3	3	3	2.46s	✓3.17s	✗4.21s	✗2.24s	✓3.54s
Tree	3	3	3	1.61s	✓2.64s	✓3.40s	✓3.08s	✓3.12s
DSL	43	43	23	2.89s	✓5.46s	✗3.85s	✗4.19s	✓6.54s
LH-FingerTree	2	1	1	5.55s	✓5.60s	✓5.57s	✓5.64s	✓5.95s
LH-T1013	1	1	1	1.11s	✓1.06s	✓1.00s	✓1.02s	✓1.06s
LH-T1025	2	2	2	1.03s	✓1.05s	✓1.08s	✓1.07s	✓1.13s
LH-T1548	1	1	1	1.45s	✓1.33s	✓1.38s	✓1.32s	✓1.45s
LH-T1660	1	1	1	1.09s	✓1.12s	✓1.12s	✓1.12s	✓1.20s
LH-MapReduce	4	3	2	14.38s	✓29.50s	✓518.91s	✓28.49s	✗Timeout

■ **Table 2** Results from simplification of proofs with rewriting. **Set-Dafny** is the set example from [34], **Set-Mono** describes a similar property. **List** and **Tree** are equational proofs from [48]. **DSL** is the program equivalence case study. The remaining proofs are from the Liquid Haskell test suite folder `tests/pos`, excluding those using only inductive or mutually inductive lemmas. **Orig.** is the number of non-inductive lemma applications in the original proof. **Cut** is the number of lemma applications that were removed by rewriting. **Rules** is the number of axioms encoded as rewrite rules. **Time (Orig.)** is verification time in seconds for the original proof. **LPQO** and **KBQO** are OCAs derived from the Lexicographic Path Ordering and Knuth-Bendix ordering respectively, and **Fuel** is an OCA allowing up to 5 rewrite applications per proof goal.

968 time. As expected, no ordering constraint algebra was able to complete all the proofs using
 969 rewriting; however, each proof could be verified with at least one of them.

970 The test cases `LH-FingerTree` and `LH-MapReduce` required manual axiom instantiations be-
 971 cause the structure of the term did not match the rewrite rule for the axiom. `LH-MapReduce`,
 972 requires proving the identity `op (f (take n is)) (mapReduce n f op (drop n is)) = f is`.
 973 An inductive lemma application generates the background equality `mapReduce n f op (drop`
 974 `n is) = f (drop n is)`, and a rewrite matching the term `op (f (take n is)) (f (drop n`
 975 `is))` must be instantiated to complete the proof. However, since the background equality
 976 is neither a rewrite rule nor an evaluation step, the necessary term `op (f (take n is)) (f`
 977 `(drop n is))` never appears. Therefore, it is necessary to manually instantiate the lemma.
 978 As future work, a limited form of E-matching [12] could be used to address this issue in the
 979 general case.

980 In conclusion, we’ve shown that extending Liquid Haskell to use REST enables rewriting
 981 functionality not subsumed by existing theorem provers, that REST is effective for axiom
 982 instantiation, and that REST can simplify equational proofs.

983 ■ 8 Related Work

984 **Theorem Provers & Rewriting** Term rewriting is an effective technique to automate theo-
 985 rem proving [25] supported by most standard theorem provers. § 7.1 compares, by examples,
 986 our technique with `Coq`, `Agda`, `Lean`, and `Isabelle/HOL`. In short, our approach is different
 987 because it uses user-specified rewrite rules to derive, in a terminating way, equalities that
 988 strengthen the SMT-decidable verification conditions generated during program verification.

989 **SMT Verification & Rewriting** Our rewrite rules could be encoded in SMT solvers as
 990 universally quantified equations and instantiated using *E-matching* [12], i.e., a common
 991 algorithm for quantifier instantiation. E-matching might generate matching loops leading

992 to unpredictable divergence. [32] refers to this unpredictable behavior of E-matching as the
993 “the butterfly effect” and partially addresses it by detecting formulas that could give rise to
994 matching loops. Our approach circumvents unpredictability by using the terminating REST
995 algorithm to instantiate the rewrite rules outside of the SMT solver.

996 Z3 [13] and CVC4 [6] are state-of-the-art SMT solvers; both support theory-specific rewrite
997 rules internally. Recent work [40] enables user-provided rewrite rules to be added to CVC4.
998 However, using the SMT solver as a rewrite engine offers little control over rewrite rule
999 instantiation, which is necessary for ensuring termination.

1000 **Rewriting in Haskell** Haskell itself has used various notions of rewriting for program veri-
1001 fication. GHC supports the `RULES` pragma with which the user can specify unchecked, quan-
1002 tified expression equalities that are used at compile time for program optimization. [10] pro-
1003 poses Inspection Testing as a way to check such rewrite rules using runtime execution and
1004 metaprogramming, while [22] prove rewrite rules via metaprogramming and user-provided
1005 hints. In a work closely related to ours, Zeno [44] is using rewriting, induction, and further
1006 heuristics to provide lemma discovery and fully automatic proof generation of inductive
1007 properties. Unlike our approach, Zeno’s syntax is restricted (e.g., it does not allow for ex-
1008 istentials) and it does not allow for user-provided hints when automation fails. HALO [51]
1009 enables Haskell verification by converting Haskell into logic and using an SMT solver to
1010 verify user-defined formulas. However, this approach relies on SMT quantifiers to encode
1011 user functions, thus the solver can diverge and verification becomes unpredictable.

1012 **Termination of Rewriting and Runtime Termination Checking** Early work on proving
1013 termination of rewriting using simplification orderings is described in [15]. More recent
1014 work involves dependency pairs [3] and applying the size-change termination principle [31]
1015 in the context of rewriting [47]. Tools like AProVE [24] and NaTT [54] can statically prove
1016 the termination of rewriting.

1017 In contrast, REST is not focused on statically proving termination of rewriting; rather
1018 it uses a well-founded ordering to ensure termination at runtime. This approach enables
1019 integration of arbitrary external oracles to produce rewrite applications, as a static analysis
1020 is not possible in principle. Furthermore, our approach enables nonterminating rewriting
1021 systems to be useful: REST will still apply certain rewrite rules to satisfy a proof obligation,
1022 even if the rewrite rules themselves cannot be statically shown to terminate.

1023 We choose to use a well-quasi-ordering [30] because it enables rewriting to terms that
1024 are not strictly decreasing in a simplification ordering. WQOs are commonly used in online
1025 termination checking [35], especially for program optimization techniques such as supercom-
1026 pilation [9].

1027 **Equality Saturation** In our implementation, REST passes equalities to the SMT environ-
1028 ment, ultimately used for *equality saturation* via an E-graph data structure [20]. Equality
1029 saturation has also been used for supercompilation[46]. REST does not currently exploit equal-
1030 ity saturation (unless indirectly via its oracle). However, as future work we might explore
1031 local usage of efficient E-graph implementations. (e.g., [53]) for caching the equivalence
1032 classes generated via rewrite applications.

1033 **Associative-Commutative Rewriting** Traditionally, enforcing a strict ordering on terms
1034 prevents the application of rewrite rules for associativity or commutativity (AC); this prob-
1035 lem motivates REST’s use of well-quasi orders. However, another solution is to omit the
1036 rules and instead perform the substitution step of rewriting modulo AC. Termination of the

1037 resulting system can be proved using an AC ordering [17]; the essential requirement is that
 1038 the ordering also respects AC, such that $t > u$ implies $t' > u'$ for all terms t' AC-equivalent
 1039 to t and u' AC-equivalent to u .

1040 REST's use of well-quasi-orderings enables AC axioms to be encoded as rewrite rules,
 1041 guaranteeing completeness if the AC-equivalence class of a term is a subset of the equivalence
 1042 class induced by the ordering. This is a significant practical benefit as it does not require
 1043 REST to identify AC symbols and treat them differently for unification.

1044 However, we note that treating AC axioms as rewrite rules can lead to an explosion in
 1045 the number of terms obtained via rewriting. As future work, it could be possible to extend
 1046 REST to support AC rewriting and unification in order to reduce the number of explicitly
 1047 instantiated terms.

1048 9 Conclusion

1049 We have presented REST, a novel approach to rewriting that uses an online termination check
 1050 that simultaneously considers entire families of term orderings via a newly introduced Or-
 1051 dering Constraint Algebra. We defined our algebra on well-quasi orderings that are more
 1052 permissive than standard simplification orderings, and demonstrated how to derive well-
 1053 quasi orderings from well-known simplification orderings. In addition, we proved correct-
 1054 ness, relative completeness, and (online) termination of our algorithm. Our REST approach
 1055 is designed, via a generic core algorithm and the pluggable abstraction of our OCAs, to be
 1056 simple to (re-)implement and adapt to different programming languages or efficient imple-
 1057 mentations of term ordering families. We demonstrated this by writing an implementation
 1058 of REST as a small Haskell library suitable for integration with existing verification tools. To
 1059 evaluate REST we used our library to extend Liquid Haskell, and showed that the resulting
 1060 system compares well with existing rewriting techniques, it can be used as an alternative
 1061 to E-matching based axiomatizations approaches without risking non-termination, and can
 1062 substantially simplify proofs requiring equational reasoning steps.

1063 References

- 1064 1 Agda Developers. *The Agda Language Reference, version 2.6.1*, 2020. Available electronically
 1065 at <https://agda.readthedocs.io/en/v2.6.1/language/index.html>.
- 1066 2 Nada Amin, K Rustan M Leino, and Tiark Rompf. Computing with an smt solver. In
 1067 *International Conference on Tests and Proofs*, pages 20–35. Springer, 2014.
- 1068 3 Thomas Arts and Jürgen Giesl. Termination of term rewriting using dependency pairs. *The-*
 1069 *oretical Computer Science*, 236(1):133–178, April 2000. URL: [http://www.sciencedirect.com/](http://www.sciencedirect.com/science/article/pii/S0304397599002078)
 1070 [science/article/pii/S0304397599002078](http://www.sciencedirect.com/science/article/pii/S0304397599002078), doi:10.1016/S0304-3975(99)00207-8.
- 1071 4 Jeremy Avigad, Leonardo de Moura, and Soonho Kong. *Theorem Proving in Lean, Release*
 1072 *3.20.0*, September 2020. p 73. URL: [https://leanprover.github.io/theorem_proving_in_lean/](https://leanprover.github.io/theorem_proving_in_lean/theorem_proving_in_lean.pdf)
 1073 [theorem_proving_in_lean.pdf](https://leanprover.github.io/theorem_proving_in_lean/theorem_proving_in_lean.pdf).
- 1074 5 Jeremy Avigad, Gabriel Ebner, and Sebastian Ullrich. *The Lean Reference Manual, Release*
 1075 *3.3.0*, 2018. URL: https://leanprover.github.io/reference/lean_reference.pdf.
- 1076 6 Clark Barrett, Christopher L. Conway, Morgan Deters, Liana Hadarean, Dejan Jovanović,
 1077 Tim King, Andrew Reynolds, and Cesare Tinelli. CVC4. In Ganesh Gopalakrishnan and
 1078 Shaz Qadeer, editors, *Proceedings of the 23rd International Conference on Computer Aided*
 1079 *Verification (CAV '11)*, volume 6806 of *Lecture Notes in Computer Science*, pages 171–177.
 1080 Springer, July 2011. Snowbird, Utah. URL: [http://www.cs.stanford.edu/~barrett/pubs/BCD+](http://www.cs.stanford.edu/~barrett/pubs/BCD+11.pdf)
 1081 [11.pdf](http://www.cs.stanford.edu/~barrett/pubs/BCD+11.pdf).

- 1082 7 N. Becker, P. Müller, and A. J. Summers. The axiom profiler: Understanding and debugging
1083 smt quantifier instantiations. In *Tools and Algorithms for the Construction and Analysis of*
1084 *Systems (TACAS) 2019*, LNCS, pages 99–116. Springer-Verlag, 2019.
- 1085 8 Maximilian Bolingbroke and Simon Peyton Jones. Supercompilation by evaluation. *SIGPLAN*
1086 *Not.*, 45(11):135146, September 2010. doi:10.1145/2088456.1863540.
- 1087 9 Maximilian Bolingbroke, Simon Peyton Jones, and Dimitrios Vytiniotis. Termination combi-
1088 nators forever. In *Proceedings of the 4th ACM symposium on Haskell*, pages 23–34, 2011.
- 1089 10 Joachim Breitner. A promise checked is a promise kept: inspection testing. In Nicolas
1090 Wu, editor, *Proceedings of the 11th ACM SIGPLAN International Symposium on Haskell,*
1091 *Haskell@ICFP 2018, St. Louis, MO, USA, September 27-17, 2018*, pages 14–25. ACM, 2018.
1092 doi:10.1145/3242744.3242748.
- 1093 11 The Coq Development Team. *The Coq Reference Manual, version 8.11.2*, 2020. Available
1094 electronically at <http://coq.inria.fr/refman>.
- 1095 12 Leonardo de Moura and Nikolaj Bjørner. Efficient e-matching for smt solvers. In Frank
1096 Pfenning, editor, *Automated Deduction – CADE-21*, pages 183–198, Berlin, Heidelberg, 2007.
1097 Springer Berlin Heidelberg.
- 1098 13 Leonardo De Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *International*
1099 *conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages
1100 337–340. Springer, 2008.
- 1101 14 Nachum Dershowitz. A note on simplification orderings. *Information Processing Letters*,
1102 9(5):212–215, 1979. doi:[https://doi.org/10.1016/0020-0190\(79\)90071-1](https://doi.org/10.1016/0020-0190(79)90071-1).
- 1103 15 Nachum Dershowitz. Orderings for term-rewriting systems. *Theoretical computer science*,
1104 17(3):279–301, 1982.
- 1105 16 Nachum Dershowitz. Termination of rewriting. *Journal of symbolic computation*, 3(1-2):69–
1106 115, 1987.
- 1107 17 Nachum Dershowitz, Jieh Hsiang, N Alan Josephson, and David A Plaisted. Associative-
1108 commutative rewriting. In *IJCAI*, pages 940–944, 1983.
- 1109 18 Nachum Dershowitz and Zohar Manna. Proving termination with multiset orderings. In Her-
1110 mann A. Maurer, editor, *Automata, Languages and Programming*, Lecture Notes in Computer
1111 Science, pages 188–202, Berlin, Heidelberg, 1979. Springer. doi:10.1007/3-540-09510-1_15.
- 1112 19 David Detlefs, Greg Nelson, and James B. Saxe. Simplify: A theorem prover for program
1113 checking. *J. ACM*, 52(3):365–473, May 2005. URL: <http://doi.acm.org/10.1145/1066100.1066102>, doi:10.1145/1066100.1066102.
- 1115 20 David Detlefs, Greg Nelson, and James B Saxe. Simplify: a theorem prover for program
1116 checking. *Journal of the ACM (JACM)*, 52(3):365–473, 2005.
- 1117 21 Claire Dross, Sylvain Conchon, Johannes Kanig, and Andrei Paskevich. Adding decision pro-
1118 cedures to smt solvers using axioms with triggers. *Journal of Automated Reasoning*, 56(4):387–
1119 457, 2016.
- 1120 22 Andrew Farmer, Neil Sculthorpe, and Andy Gill. Reasoning with the hermit: Tool support
1121 for equational reasoning on ghc core programs. In *Proceedings of the 2015 ACM SIGPLAN*
1122 *Symposium on Haskell*, Haskell 15, page 2334, New York, NY, USA, 2015. Association for
1123 Computing Machinery. doi:10.1145/2804302.2804303.
- 1124 23 Jean-Christophe Filliâtre and Andrei Paskevich. Why3—where programs meet provers. In
1125 Matthias Felleisen and Philippa Gardner, editors, *Programming Languages and Systems*
1126 *(ESOP)*, volume 7792 of *Lecture Notes in Computer Science*, pages 125–128. Springer, 2013.
- 1127 24 Jürgen Giesl, Cornelius Aschermann, Marc Brockschmidt, Fabian Emmes, Florian Frohn,
1128 Carsten Fuhs, Jera Hensel, Carsten Otto, Martin Plücker, Peter Schneider-Kamp, et al. Ana-
1129 lyzing program termination and complexity automatically with approve. *Journal of Automated*
1130 *Reasoning*, 58(1):3–31, 2017.
- 1131 25 Jieh Hsiang, Hélène Kirchner, Pierre Lescanne, and Michaël Rusinowitch. The term rewriting
1132 approach to automated theorem proving. *The Journal of Logic Programming*, 14(1):71–99,

- 1133 October 1992. URL: <http://www.sciencedirect.com/science/article/pii/0743106692900477>,
1134 doi:10.1016/0743-1066(92)90047-7.
- 1135 26 Gerard Huet. Confluent reductions: Abstract properties and applications to term rewriting
1136 systems. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*,
1137 SFCS '77, page 3045, USA, 1977. IEEE Computer Society. doi:10.1109/SFCS.1977.9.
- 1138 27 Stéphane Kaplan. Conditional rewrite rules. *Theoretical Computer Science*, 33(2):175–
1139 193, 1984. URL: <https://www.sciencedirect.com/science/article/pii/0304397584900872>, doi:
1140 [https://doi.org/10.1016/0304-3975\(84\)90087-2](https://doi.org/10.1016/0304-3975(84)90087-2).
- 1141 28 J. W. Klop. *Term Rewriting Systems*, page 1116. Oxford University Press, Inc., USA, 1993.
- 1142 29 Donald E Knuth and Peter B Bendix. Simple word problems in universal algebras. In
1143 *Automation of Reasoning*, pages 342–376. Springer, 1983.
- 1144 30 Joseph B Kruskal. The theory of well-quasi-ordering: A frequently discovered concept.
1145 *Journal of Combinatorial Theory, Series A*, 13(3):297–305, November 1972. URL: <http://www.sciencedirect.com/science/article/pii/0097316572900635>, doi:10.1016/0097-3165(72)
1146 90063-5.
- 1148 31 Chin Soon Lee, Neil D. Jones, and Amir M. Ben-Amram. The size-change principle for
1149 program termination. In *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on*
1150 *Principles of Programming Languages*, POPL 01, page 8192, New York, NY, USA, 2001.
1151 Association for Computing Machinery. doi:10.1145/360204.360210.
- 1152 32 K. R. M. Leino and Clément Pit-Claudel. Trigger Selection Strategies to Stabilize Program
1153 Verifiers. In Swarat Chaudhuri and Azadeh Farzan, editors, *Computer Aided Verification*,
1154 *Lecture Notes in Computer Science*, pages 361–381, Cham, 2016. Springer International Pub-
1155 lishing. doi:10.1007/978-3-319-41528-4_20.
- 1156 33 K. Rustan M. Leino. Dafny: An automatic program verifier for functional correctness. In
1157 *Proceedings of the 16th International Conference on Logic for Programming, Artificial Intel-*
1158 *ligence, and Reasoning*, LPAR10, page 348370, Berlin, Heidelberg, 2010. Springer-Verlag.
- 1159 34 K Rustan M Leino and Nadia Polikarpova. Verified calculations. In *Working Conference on*
1160 *Verified Software: Theories, Tools, and Experiments*, pages 170–190. Springer, 2013.
- 1161 35 Michael Leuschel. Homeomorphic Embedding for Online Termination of Symbolic Meth-
1162 ods. In Gerhard Goos, Juris Hartmanis, Jan van Leeuwen, Torben Æ. Mogensen, David A.
1163 Schmidt, and I. Hal Sudborough, editors, *The Essence of Computation*, volume 2566,
1164 pages 379–403. Springer Berlin Heidelberg, Berlin, Heidelberg, 2002. Series Title: Lec-
1165 ture Notes in Computer Science. URL: http://link.springer.com/10.1007/3-540-36377-7_17,
1166 doi:10.1007/3-540-36377-7_17.
- 1167 36 Jia Meng and Lawrence C Paulson. Translating higher-order clauses to first-order clauses.
1168 *Journal of Automated Reasoning*, 40(1):35–60, 2008.
- 1169 37 P. Müller, M. Schwerhoff, and A. J. Summers. Viper: A verification infrastructure for
1170 permission-based reasoning. In *VMCAI*, 2016.
- 1171 38 Tobias Nipkow, Markus Wenzel, and Lawrence C. Paulson. *Isabelle/HOL: A Proof Assistant*
1172 *for Higher-Order Logic*. Springer-Verlag, 2020.
- 1173 39 Ulf Norell. Dependently typed programming in agda. In *Proceedings of the 6th International*
1174 *Conference on Advanced Functional Programming*, AFP08, page 230266, Berlin, Heidelberg,
1175 2008. Springer-Verlag.
- 1176 40 Andres Nötzli, Andrew Reynolds, Haniel Barbosa, Aina Niemetz, Mathias Preiner, Clark
1177 Barrett, and Cesare Tinelli. Syntax-guided rewrite rule enumeration for smt solvers. In
1178 Mikoláš Janota and Inês Lynce, editors, *Theory and Applications of Satisfiability Testing –*
1179 *SAT 2019*, pages 279–297, Cham, 2019. Springer International Publishing.
- 1180 41 Lawrence C Paulson and Kong Woei Susanto. Source-level proof reconstruction for interactive
1181 theorem proving. In *International Conference on Theorem Proving in Higher Order Logics*,
1182 pages 232–245. Springer, 2007.
- 1183 42 Lawrence C Paulsson and Jasmin C Blanchette. Three years of experience with sledgehammer,
1184 a practical link between automatic and interactive theorem provers. In *Proceedings of the 8th*

- 1185 *International Workshop on the Implementation of Logics (IWIL-2010)*, Yogyakarta, Indonesia.
1186 *EPiC*, volume 2, 2012.
- 1187 43 Julien Signoles, Pascal Cuoq, Florent Kirchner, Nikolai Kosmatov, Virgile Prevosto, and
1188 Boris Yakobowski. Frama-c: a software analysis perspective. volume 27, 10 2012. doi:
1189 10.1007/s00165-014-0326-7.
- 1190 44 William Sonnex, Sophia Drossopoulou, and Susan Eisenbach. Zeno: An automated prover
1191 for properties of recursive data structures. In Cormac Flanagan and Barbara König, editors,
1192 *Tools and Algorithms for the Construction and Analysis of Systems*, pages 407–421, Berlin,
1193 Heidelberg, 2012. Springer Berlin Heidelberg.
- 1194 45 Nikhil Swamy, Catalin Hritcu, Chantal Keller, Aseem Rastogi, Antoine Delignat-Lavaud, Si-
1195 mon Forest, Karthikeyan Bhargavan, Cédric Fournet, Pierre-Yves Strub, Markulf Kohlweiss,
1196 Jean-Karim Zinzindohoué, and Santiago Zanella-Béguelin. Dependent types and multi-
1197 monadic effects in F*. In *43rd ACM SIGPLAN-SIGACT Symposium on Principles of*
1198 *Programming Languages (POPL)*, pages 256–270. ACM, January 2016. URL: <https://www.fstar-lang.org/papers/mumon/>.
- 1200 46 Ross Tate, Michael Stepp, Zachary Tatlock, and Sorin Lerner. Equality saturation: a new
1201 approach to optimization. In *POPL '09: Proceedings of the 36th annual ACM SIGPLAN-*
1202 *SIGACT symposium on Principles of Programming Languages*, pages 264–276, New York,
1203 NY, USA, 2009. ACM. URL: <http://www.cs.cornell.edu/~ross/publications/eqsat/>, doi:
1204 <http://doi.acm.org/10.1145/1480881.1480915>.
- 1205 47 René Thiemann and Jürgen Giesl. Size-Change Termination for Term Rewriting. volume
1206 2706, pages 264–278, March 2007. doi:10.1007/3-540-44881-0_19.
- 1207 48 Niki Vazou, Joachim Breitner, Rose Kunkel, David Van Horn, and Graham Hutton. Theorem
1208 proving for all: equational reasoning in liquid Haskell (functional pearl). In *Proceedings of*
1209 *the 11th ACM SIGPLAN International Symposium on Haskell*, Haskell 2018, pages 132–144,
1210 St. Louis, MO, USA, September 2018. Association for Computing Machinery. doi:10.1145/
1211 3242744.3242756.
- 1212 49 Niki Vazou, Eric L. Seidel, Ranjit Jhala, Dimitrios Vytiniotis, and Simon Peyton-Jones. Re-
1213 finement types for haskell. In *Proceedings of the 19th ACM SIGPLAN International Confer-*
1214 *ence on Functional Programming*, ICFP 14, page 269282, New York, NY, USA, 2014. Associ-
1215 ation for Computing Machinery. doi:10.1145/2628136.2628161.
- 1216 50 Niki Vazou, Anish Tondwalkar, Vikraman Choudhury, Ryan G. Scott, Ryan R. Newton, Philip
1217 Wadler, and Ranjit Jhala. Refinement reflection: Complete verification with smt. *Proc. ACM*
1218 *Program. Lang.*, 2(POPL), December 2017. doi:10.1145/3158141.
- 1219 51 Dimitrios Vytiniotis, Simon Peyton Jones, Koen Claessen, and Dan Rosén. Halo: Haskell
1220 to logic through denotational semantics. In *Proceedings of the 40th annual ACM SIGPLAN-*
1221 *SIGACT symposium on Principles of programming languages*, pages 431–442, 2013.
- 1222 52 Philip Wadler. Deforestation: transforming programs to eliminate trees. *Theoretical Computer*
1223 *Science*, 73(2):231 – 248, 1990. URL: [http://www.sciencedirect.com/science/article/pii/](http://www.sciencedirect.com/science/article/pii/030439759090147A)
1224 [030439759090147A](http://www.sciencedirect.com/science/article/pii/030439759090147A), doi:[https://doi.org/10.1016/0304-3975\(90\)90147-A](https://doi.org/10.1016/0304-3975(90)90147-A).
- 1225 53 Max Willsey, Chandrakana Nandi, Yisu Remy Wang, Oliver Flatt, Zachary Tatlock, and
1226 Pavel Panchekha. Egg: Fast and extensible equality saturation. *Proceedings of the ACM on*
1227 *Programming Languages*, 5(POPL):1–29, 2021.
- 1228 54 Akihisa Yamada, Keiichirou Kusakari, and Toshiki Sakabe. Nagoya termination tool. In
1229 *Rewriting and Typed Lambda Calculi*, pages 466–475. Springer, 2014.

1230 **A Metaproperty Proofs**

1231 ► **REST Invariant 1** (Path Invariant). *For any execution of $REST(\mathcal{A}, R, t_0, \mathcal{E})$, at the start of*
 1232 *any iteration of the main loop, for each $(ts, c) \in p$, the list ts is a path of $R + \mathcal{E}$ starting*
 1233 *from t_0 .*

1234 **Proof.** By induction on the loop iterations of the algorithm. p is initialized with the single
 1235 element $([t_0], c)$. $[t_0]$ is a valid path of $R + \mathcal{E}$, because it only contains a single term; clearly
 1236 this path also starts with t_0 .

1237 At each loop iteration, new elements are potentially pushed to p . Suppose the path ts is
 1238 popped from p at the beginning of the loop. The element to be pushed is a pair $(ts ++ [t'], c)$
 1239 where $last(ts) \rightarrow_{R+\mathcal{E}} t'$. This exactly satisfies the inductive hypothesis: if ts is a path of
 1240 $R + \mathcal{E}$, then $ts ++ [t']$ is also a path of $R + \mathcal{E}$. Furthermore, this operation preserves the head
 1241 of the list: t_0 is still the first element. ◀

1242 ► **Theorem 18** (Completeness w.r.t. \mathcal{E}). *For all R, u , and t_0 , if $t_0 \rightarrow_{\mathcal{E}}^* u$, then $u \in$
 1243 $REST(\mathcal{A}, R, t_0, \mathcal{E})$.*

1244 **Proof.** The proof goes by induction on the number of steps of the path.

1245 Assume the path has n steps: $t_0 \rightarrow_{\mathcal{E}} t_1 \rightarrow_{\mathcal{E}} \dots \rightarrow_{\mathcal{E}} t_{n-1} \rightarrow_{\mathcal{E}} t_n \equiv u$.

1246 For the base case, $n = 0$ and $u \equiv t_0$. Since p is initialized with $([t_0], \top)$, by the Invariant 2,
 1247 $t \in REST(\mathcal{A}, R, t_0, \mathcal{E})$.

1248 For the inductive case, assume that $t_0 \rightarrow_{\mathcal{E}}^* t_{n-1} \rightarrow_{\mathcal{E}} t_n$. By inductive hypothesis, $t_{n-1} \in$
 1249 $REST(\mathcal{A}, R, t_0, \mathcal{E})$. When t_{n-1} was added in the result, it was the last element of a path ts
 1250 that was popped from the stack p . Since $t_{n-1} \rightarrow_{\mathcal{E}} t_n$, we split cases on whether or not
 1251 $t_n \in ts$. If $t_n \in ts$, then by Invariant 2 $t_n \in REST(\mathcal{A}, R, t_0, \mathcal{E})$. Otherwise, $(ts ++ [t_n], c)$ will
 1252 be pushed into p and, again, by Invariant 2 it will appear in the output. ◀

1253 ► **Theorem 19** (Relative Completeness). *For all R, u , and t_0 , if $t_0 \rightarrow_{R+\mathcal{E}}^* u$ and there exists*
 1254 *an ordering $\succ \in \gamma(\top)$ that orients the path justifying $t_0 \rightarrow_{R+\mathcal{E}}^* u$, then $u \in REST(\mathcal{A}, R, t_0, \mathcal{E})$.*

1255 First, we observe the (somewhat standard) property that if any path justifies $t_0 \rightarrow_{R+\mathcal{E}}^* u$,
 1256 there is a *duplicate-free variant* of such path (intuitively, obtained by cutting out all subpaths
 1257 leading from a term to itself).

1258 Below, we prove that if $t_0 \rightarrow_{R+\mathcal{E}}^* u$ and the ordering \succ orients the path, then a duplicate-
 1259 free variant path ts belongs in the stack p with some constraints c and $\succ \in \gamma(c)$.

1260 ► **REST Invariant 3.** *For any execution of $REST(\mathcal{A}, R, t_0, \mathcal{E})$, if $t_0 \rightarrow_{R+\mathcal{E}}^* t_n$ and $\succ \in \gamma(\top)$ is*
 1261 *an ordering that orients $t_0 \rightarrow_{R+\mathcal{E}}^* u$, then at some iteration of the main loop, a duplicate-free*
 1262 *variant path ts of this path is stored in p , with some ordering constraints c and $\succ \in \gamma(c)$.*

1263 **Proof.** The proof goes by strong induction on the length $n+1$ of the path justifying $t_0 \rightarrow_{R+\mathcal{E}}^*$
 1264 t_n .

1265 First, consider the case $n = 0$, where the path is $[t_0]$ and the constraints \top . $([t_0], \top) \in p$
 1266 by initialization and trivially $\succ \in \gamma(\top)$.

1267 Otherwise, when $n > 0$, assume that $t_0 \rightarrow_{R+\mathcal{E}}^* t_{n-1} \rightarrow_{R+\mathcal{E}} t_n$. If there are any duplicate
 1268 terms in this path, a duplicate-free variant exists of shorter length, and we can conclude by
 1269 our induction hypothesis. Otherwise, consider this path with the last element t_n removed.
 1270 Being already duplicate-free, by our induction hypothesis we must have that, at some iter-
 1271 ation of our main loop, this path is contained in p along with a constraint c_{n-1} such that
 1272 $\succ \in \gamma(c_{n-1})$. By the assumption that \succ orients the original path, in particular we must
 1273 have $t_{n-1} \succ t_n$, and so, by Def. 14, $\succ \in \gamma(refine(c_{n-1}, t, t'))$ and therefore $refine(c_{n-1}, t, t')$

1274 is satisfiable. Therefore, the original path will be pushed to p with this constraint in this
 1275 loop iteration. ◀

1276 **Proof.** The proof is similar to Theorem 18, but now we need to also show that the relation
 1277 that orients the path satisfies all the ordering constraints generated by the respective REST
 1278 path. By Invariant 3, at some iteration of the main loop, there must be some path ending
 1279 in u contained in p . Then, by Invariant 2 it follows that all the elements of the path, thus
 1280 also u , belong in the result.

1281 ◀

1282 ▶ **Theorem 22** (Termination of REST). *For any finite set of rewriting rules R , if:*

- 1283 1. $\rightarrow_{\mathcal{E}}$ is normalizing and bounded,
- 1284 2. The refine and sat functions from \mathcal{A} are decidable (always-terminating, in an implemen-
 1285 tation),
- 1286 3. \mathcal{A} is well-founded,

1287 then, for all terms t_0 , $\text{REST}(\mathcal{A}, R, t_0, \mathcal{E})$ terminates.

1288 **Proof.** At every iteration of REST, a path with length n is popped off the stack and due
 1289 to Requirement 1, and the fact that only a finite number of new terms can be generated
 1290 by single applications of the rules R to an arbitrary term, a finite number of paths with
 1291 length $n + 1$ is pushed on. Therefore, REST implicitly builds (via its set of paths p) a *finitely-*
 1292 *branching* tree starting from t_0 . For REST to not terminate, there must be an infinite path
 1293 down the tree (note that Requirement 2 eliminates the possibility that the operations called
 1294 from the ordering constraint algebra cause non-termination).

1295 Consider an arbitrary path down the tree explored by REST, represented by the (ts, c)
 1296 pairs iteratively generated. Firstly, due to the first condition in the foreach of REST (cf.
 1297 Figure 3), this path will remain duplicate-free. By Requirement 3, at only finitely many
 1298 steps is the constraint tracked *strictly* refined. Consider then, the postfix of the path after
 1299 the last time that this happens; at every step, the constraint c remains identical. The
 1300 normalization assumption (Requirement 1) of \mathcal{E} entails that this path contains no infinite
 1301 sequence of steps all justified by \mathcal{E} . However, for each step justified instead by a rewriting
 1302 step from R , the additional condition $\text{sat}(c)$ must hold; by Def. 14 this means that there is
 1303 some $\succ \in \gamma(c)$ which orients all of these steps. As \succ must be an instance of a term ordering
 1304 family (Def. 4), it is a thin well-quasi-order. Therefore \succ can only orient a finite number of
 1305 steps, and the path down the tree must be finite.

1306 Since every path in the finitely-branching tree explored is finite, the algorithm (always)
 1307 terminates. ◀

1308 **B Proofs on Orderings**

1309 ► **Lemma 23.** *If $T \succ_{M(X)} U$, then $T \succ_{M(X)} U'$ for all $U' \subset U$.*

1310 **Proof.** It is sufficient to show that $T \succ_{M(X)} U$ implies $T \succ_{M(X)} (U - u')$, for any $u' \in U$,
 1311 since the subset can be obtained by removing a finite number of elements. That is, if U' was
 1312 obtained by removing elements u_1, \dots, u_n from U , we can show that $T \succ_{M(X)} (U \setminus \{u_1\})$
 1313 implies $T \succ_{M(X)} (U \setminus \{u_1, u_2\})$ and so on.

1314 The proof goes by induction on the size of T and case analysis on $T \succ_{M(X)} U$.

1315 For case one there are no u' in U , so the proof holds vacuously.

1316 For case two, we have either $u = u'$ or $u \neq u'$. If $u = u'$, a proof of $T \succ_{M(X)} (U - u)$
 1317 can be made by modifying the proof of $(T - t) \succ_{M(X)} (U - u)$. The base case of that proof
 1318 must be of the form $T' \succ_{M(X)} \emptyset$. We modify the base case to be $(T' + t) \succ_{M(X)} \emptyset$. Each
 1319 recursive case is also modified to replace T' with $(T' + t)$, yielding $(T' + t) \succ_{M(X)} (U - u)$
 1320 $= T \succ_{M(X)} (U - u)$, as required. The proof that $T \succ_{M(X)} (U - u')$ for all other $u' \in U$ is
 1321 obtained by induction. By the inductive hypothesis, we have $(T - t) \succ_{M(X)} (U - u - u')$, since
 1322 $u \neq u'$, we also have $u \in (U - u')$. Therefore applying case two we get $T \succ_{M(X)} (U - u')$.

1323 For case three, we have either $u' < t$ or $u' \not< t$. If $u' < t$, then the proof $(T - t) \succ_{M(X)}$
 1324 $(U \setminus \{u \in U \mid u < t\})$ is also a proof of $(T - t) \succ_{M(X)} ((U - u') \setminus \{u \in U \mid u < t\})$,
 1325 thus we obtain obtain the proof directly. The proof for all other $u' \in U$ is obtained by
 1326 induction. By the inductive hypothesis we have $(T - t) \succ_{M(X)} ((U \setminus \{u \in U \mid u < t\}) - u')$.
 1327 Then, applying the same top-level proof yields $T \succ_{M(X)} (U - u')$, since u' is not in the set
 1328 $\{u \in U \mid u < t\}$.

1329

◀

1330 ► **Lemma 24.** *If \succ_X is a quasi-order, then the multiset extension $\succ_{M(X)}$ is also a quasi-*
 1331 *order.*

1332 **Proof.** To show that $\succ_{M(X)}$ is a quasi-order, we define a single-step version \succ_{mul} , and show
 1333 that $T \succ_{M(X)} U$ if and only if $T \succ_{mul^*} U$, where \succ_{mul^*} is the reflexive transitive closure of
 1334 \succ_{mul} .

1335 We define \succ_{mul} as:

- 1336 1. For all elements t, u if $t \in T$ and $u \approx t$, then $T \succ_{mul} (T - t + u)$
- 1337 2. For all elements $t \in T$ and finite multisets U , if $t > u$ for all $u \in U$, then $T \succ_{mul}$
 1338 $((T - t) \cup U)$

1339 First, observe that \succ_{mul^*} is monotonic with respect to multiset union: for all multisets
 1340 T, U , and V , $T \succ_{mul^*} U$ implies $(T \cup V) \succ_{mul^*} (U \cup V)$.

1341 The reflexive case is given by $T \cup V = T \cup V$; we show the transitive case by showing
 1342 there is a correspondence for each single-step. The proof for each case assumes an arbitrary
 1343 multiset V .

1344 In case one we must show for all $t, u \in T$, $T \succ_{mul^*} (T - t + u)$ implies $(T \cup V) \succ_{mul^*}$
 1345 $((T - t + u) \cup V)$. t and u are also in $T \cup V$, therefore we have $(T \cup V) \succ_{mul^*} ((T \cup V) - t + u)$.
 1346 We have $(T \cup V) - t + u = (T - t + u) \cup V$, giving us the desired result. Case two is similar:
 1347 $t \in T$ implies $t \in (T \cup V)$, and $((T \cup V) - t) \cup U = ((T - t) \cup U) \cup V$ for all U, V .

1348 Now we show the if direction by case analysis.

1349 Case 1: $U = \emptyset$.

1350 If $T = \emptyset$, then we have $T \succ_{mul^*} U$ via reflexivity. Otherwise we can select an arbitrary t

1351 to remove from T , and by definition of \geq_{mul} we have $T \geq_{mul} ((T-t) \cup \emptyset)$. Then by in-
 1352 duction on the size of T we have $((T-t) \cup \emptyset) \geq_{mul*} \emptyset$. Then $T \geq_{mul} ((T-t) \cup \emptyset) \geq_{mul*} \emptyset$,
 1353 as required.

1354
 1355 Case 2: $t \in T \wedge u \in U \wedge t \approx u \wedge (T-t) \geq_{mul} (U-u)$.

1356 Let $T' = T-t$ and $U' = U-u$. Then we have $(T'+t) \geq_{mul} (T'+u)$ by definition and
 1357 $T' \succ_{M(X)} U'$ implies $T'+u \geq_{mul*} U'+u$ via the inductive hypothesis and monotonicity.
 1358 Thus $T = (T'+t) \geq_{mul} (T'+u) \geq_{mul*} (U'+u) = U$ as required.

1359
 1360 Case 3: $t \in T \wedge (T-t) \geq_{mul} (U \setminus \{u \in U \mid u < t\})$

1361 Partition U into two sets U_1 and U_2 where $U_1 = \{u \in U \mid u \not< t\}$ and $U_2 = \{u \in$
 1362 $U \mid u < t\}$. By definition we have $U = U_1 \cup U_2$. As before $T' = T-t$. Then we have
 1363 $(T'+t) \geq_{mul} (T' \cup U_2)$. $T' \succ_{M(X)} U_1$ implies $(T' \cup U_2) \geq_{mul*} (U_1 \cup U_2)$ via monotonicity
 1364 and induction. Thus $T = (T'+t) \geq_{mul} (T' \cup U_2) \geq_{mul*} (U_1 \cup U_2) = U$ as required.

1365 Now the only-if direction. First we have that $\succ_{M(X)}$ is reflexive via induction on size
 1366 with base case $T = U = \emptyset$ handled by case 1, and recursive case by case 2, similar to above,
 1367 we remove an arbitrary t from T . Now we show how to handle one or more steps from \geq_{mul}
 1368 in a single step of $\succ_{M(X)}$.

1369 The key observation is that all elements u of U , have exactly one “responsible” element
 1370 t in T that justifies $T \geq_{mul*} U$: we must have either $t > u$ or $t \approx u$ (in which case t is
 1371 uniquely responsible for u and no other elements of U). To prove $T \succ_{M(X)} U$, for each t in
 1372 T , we recursively build a tuple (T', U', p) where T' , and U' are multisets and p is the proof
 1373 that $T' \succ_{M(X)} U'$. The tuple is initialized to $(\emptyset, \emptyset, U = \emptyset)$.

1374 For each t uniquely responsible for one u , we update the tuple to $(T'+t, U'+u, t \in$
 1375 $T \wedge u \in U \wedge t \approx u \wedge p)$. The new proof state is valid because by induction we have p being
 1376 a proof of $T' \succ_{M(X)} U'$, as required.

1377 Now consider each $t \in T$ where t justified some multiset U'' . By induction, we have a
 1378 proof of $T' \succ_{M(X)} U'$; we need a proof that $T' \succ_{M(X)} ((U' \cup U'') \setminus \{u \in (U' \cup U'') \mid u < t\})$.
 1379 Since we have $t > u$ for all $u \in U''$, this simplifies to: $T' \succ_{M(X)} (U' \setminus \{u \in U' \mid u < t\})$,
 1380 which we can obtain via the hypothesis $T' \succ_{M(X)} U'$ and lemma 23.

1381 ◀

1382 ▶ **Lemma 25.** *If \succ_X is a well-quasi-order, the strict part of its multiset extension defined*
 1383 *as $t >_{M(X)} u$ if $t \succ_{M(X)} u$ and $u \not\prec_{M(X)} t$ is a well-founded order.*

1384 **Proof.** This proof operates on the single-step relation defined in 24. Proving the well-
 1385 founded property is done by showing that an infinite descent in $>_{M(X)}$ would correspond
 1386 to an infinite descent in the underlying ordering.

1387 Now, consider a tree built from an infinite path T_1, T_2, \dots of multisets related by $\succ_{M(X)}$.
 1388 With the exception of special nodes \top and \perp , each node in the tree represents an element in
 1389 a multiset, and the vertices connect the elements to the smaller ones they were replaced with
 1390 via an application of \geq_{mul} . Crucially, every edge represents an descent in a well-founded
 1391 order.

1392 The tree is constructed as follows: let \top be the root of the tree, and let the elements
 1393 of T_1 be the children of \top . Then, for each T_i in the infinite list, it was either obtained by
 1394 replacing some element in T_{i-1} with a same-sized element, or by removing some element t
 1395 and replacing it with a finite number of smaller elements ts .

1396 In the former case, the tree is not modified.

1397 In the latter case, if $ts = \emptyset$, add a single child \perp to the t in the tree. Otherwise, let ts be
 1398 the children of t .

1399 Now, we note that the case one of \succ_{mul} is symmetric. Therefore, each pair of terms
 1400 related by $>_{M(X)}$ must correspond to at least one step in case two of \succ_{mul} , Therefore in an
 1401 infinite path of terms related by $>_{M(X)}$ contains an infinite number of applications of case
 1402 two in \succ_{mul} .

1403 Therefore, an infinite number of vertices will be added to the tree. Since the tree is finitely
 1404 branching, it must have an infinitely descending path. However, this infinitely descending
 1405 path would correspond to an infinite descent in the underlying ordering, contradicting that
 1406 hypothesis that \succ_X is a WQO. \blacktriangleleft

1407 **► Lemma 26.** *If $\succ_{\mathcal{F}}$ is a total quasi-ordering, then \succ_{rpo} is a quasi-simplification ordering.*

1408 **Proof.** We must show that \succ_{rpo} is a quasi-ordering, i.e it is reflexive and transitive; and
 1409 also that it satisfies the replacement, subterm, and deletion properties.

1410 Reflexivity occurs via case 3 and 24. Replacement and deletion follow from case 3 of RPO
 1411 and the definition of the multiset ordering.

1412 To prove the subterm property, we show a slightly stronger property: for all terms
 1413 $t = f(t_1, \dots, t_m)$ and (not necessarily immediate) subterms $u = g(u_1, \dots, u_n)$, $t >_{rpo} u$.
 1414 The proof goes by induction on the term size, where terms are bigger than their subterms,
 1415 and by case analysis on the relationship between f and g . Because $\succ_{\mathcal{F}}$ is total, we have
 1416 either $f >_{\mathcal{F}} g$, $f \approx g$, or $g >_{\mathcal{F}} f$.

1417 If $f >_{\mathcal{F}} g$, then to get $t \succ_{rpo} u$ we must show $\{t\} >_{M(rpo)} \{u_1, \dots, u_n\}$. Via induction,
 1418 we have $t >_{rpo} u_i$ for all $1 \leq i \leq n$, as each u_i is a subterm of u . To show $u \not\succeq_{rpo} t$, observe
 1419 that we need $\{u_1, \dots, u_n\} \succ_{M(rpo)} \{t\}$. This is impossible via the inductive hypothesis and
 1420 the definition of $\succ_{M(rpo)}$: we already have $t >_{rpo} u_i$ for all u_i .

1421 If $f \approx g$, then we must show $\{t_1, \dots, t_m\} >_{M(rpo)} \{u_1, \dots, u_n\}$. If u is a direct subterm
 1422 of t , then $u = t_i$ for some i . By the inductive hypothesis we have $t_i \approx u >_{rpo} u_j$ for all
 1423 u_j , which implies $\{t_1, \dots, t_m\} >_{M(rpo)} \{u_1, \dots, u_n\}$. If u is a nested subterm, then we have
 1424 some $t_i >_{rpo} u_j$ for all u_j via the induction hypothesis: all u_j are subterms of t_i .

1425 If $g >_{\mathcal{F}} f$, to get $t \succ_{rpo} u$ then we must show $\{t_1, \dots, t_m\} \succ_{M(rpo)} \{u\}$. If u was a
 1426 direct subterm, then $t_i = u$ gives us the desired result; otherwise we have $t_i >_{rpo} u$ via
 1427 the inductive hypothesis. To show $u \not\succeq_{rpo} t$, observe that showing $u \succ_{rpo} t$ would require
 1428 $\{u\} >_{M(rpo)} \{t_1, \dots, t_m\}$. However we already have some $t_i \approx u$, which prevents this
 1429 possibility.

1430 Transitivity is also proven via induction on size. Assume we have $s = f(s_1, \dots, s_m) \succ_{rpo}$
 1431 $t = g(t_1, \dots, t_n)$ and $t \succ_{rpo} u = h(u_1, \dots, u_p)$. We proceed to show $s \succ_{rpo} u$ by for each
 1432 relationship between f , g , and h .

- 1433 1. $f >_{\mathcal{F}} g >_{\mathcal{F}} h$, or $f >_{\mathcal{F}} g > h$: Via transitivity of $>_{\mathcal{F}}$ we have $f >_{\mathcal{F}} h$, therefore we must
 1434 show $\{s\} >_{M(rpo)} \{u_1, \dots, u_p\}$. $\{s\} \succ_{M(rpo)} \{t\}$ follows from our assumption $s \succ_{rpo} t$,
 1435 and $\{t\} >_{M(rpo)} \{u_1, \dots, u_p\}$ follows from $t \succ_{rpo} u$. By the inductive hypothesis, we
 1436 have $s \succ_{rpo} t \succ_{rpo} u_i$ for all u_i , and therefore $\{s\} \succ_{M(rpo)} \{t\} >_{M(rpo)} \{u_1, \dots, u_p\}$.
- 1437 2. $h >_{\mathcal{F}} g$: There must exist some subterm t_i such that $t_i \succ_{rpo} u$. Therefore we have
 1438 $s \succ_{rpo} t_i$ and $t_i \succ_{rpo} u$, the inductive hypothesis gives us $s \succ_{rpo} t_i \succ_{rpo} u$.
- 1439 3. $g >_{\mathcal{F}} f$: There must exist some subterm s_i such that $s_i \succ_{rpo} t$. As above, using the
 1440 induction hypothesis allows us to show $s_i \succ_{rpo} u$, by the subterm property we have
 1441 $s \succ_{rpo} s_i$. We show $s \succ_{rpo} u$ by the definition of \succ_{rpo} .
- 1442 4. $f \approx g \approx h$. We clearly have $f \approx h$, we need to show $\{s_1, \dots, s_m\} \succ_{M(rpo)} \{u_1, \dots, u_p\}$,
 1443 which we have via 24.

1444

1445 ► **Theorem 27.** *If $\succ_{\mathcal{F}}$ is a total WQO, then \succ_{rpo} is a WQO.*

1446 **Proof.** To show that \succ_{rpo} is WQO, via the well-foundedness theorem of Dershowitz [15],
 1447 which states that a quasi-simplification ordering \succ' is WQO if there exists a well-quasi
 1448 ordering \succ such that $f \succ g$ implies $f(t_1, \dots, t_n) \succ' g(t_1, \dots, t_n)$.

1449 By 26 we have that \succ_{rpo} is a quasi-simplification ordering, and there exists an ordering
 1450 over function symbols to satisfy the condition of the well-foundedness theorem: namely the
 1451 underlying order $\succ_{\mathcal{F}}$ from which \succ_{rpo} is constructed.

1452

1453 ► **Theorem 28.** *If $\succ_{\mathcal{F}}$ is a total WQO, then \succ_{rpo} is thin*

1454 **Proof.** We show that for any term $t = f(t_1, \dots, t_m)$, the set of terms $\{u \mid t \approx u =$
 1455 $g(u_1, \dots, u_m)\}$ is finite.

1456 If $t \approx u$, then we must have $t \succ_{rpo} u$ and $u \succ_{rpo} t$. Assume we have $t \succ_{rpo} u$.

1457 First, we show that if $f > g$ then $u \not\succeq_{rpo} t$. Assume $u \succ_{rpo} t$, then there must have
 1458 some u_i such that $u_i \succ_{rpo} t$. But via the subterm property, we have $u >_{rpo} u_i \succ_{rpo} t$,
 1459 contradicting $t \succ_{rpo} u$.

1460 Likewise, if $g > f$, then there is some $t_i \succ_{rpo} u$. Then $t >_{rpo} t_i \succ_{rpo} u$. Therefore we
 1461 also have $u \not\succeq_{rpo} t$.

1462 Therefore, $t \approx u$ only if $f \approx g$. Since there are only a finite number of function symbols,
 1463 then to show thinness we must show that only a finite number of multisets $\{u_1, \dots, u_n\}$
 1464 such that $\{t_1, \dots, t_m\} \succ_{M(rpo)} \{u_1, \dots, u_n\}$ and $\{u_1, \dots, u_n\} \succ_{M(rpo)} \{t_1, \dots, t_m\}$. If
 1465 $\{t_1, \dots, t_m\} = \emptyset$, then the only such set is \emptyset . Otherwise, only such multisets are those where
 1466 $\{u_1, \dots, u_n\}$ is obtained from $\{t_1, \dots, t_m\}$ by removing zero or more terms t_i and replacing
 1467 them the same number of terms u_j where $t_i \approx u_j$. If $\{t_1, \dots, t_m\} \succ_{M(rpo)} \{u_1, \dots, u_n\}$ was
 1468 justified by removing t_i from $\{t_1, \dots, t_m\}$ and removing smaller terms $\{u' \mid u' < t_i\}$ from
 1469 $\{u_1, \dots, u_n\}$, then we would have $\{t_1, \dots, t_m\} >_{M(rpo)} \{u_1, \dots, u_n\}$: this corresponds to the
 1470 irreflexive single-step operation shown to form a well-founded order in lemma 25.

1471 Since the multisets contain a finite number of elements, and each term only has a finite
 1472 number of equivalent terms (by induction on term size), there are only a finite number of
 1473 such multisets.



1474
1475

C Basic Equalities and Proved Theorems in the Program Equivalence Case Study

	Name	Formula
1.	addDist	$(x * y) + (z * y) = (x + z) * y$
2.	subDist	$(x * y) - (z * y) = (x - z) * y$
3.	times2Plus	$x * 2 = x + x$
4.	plus0	$x + 0 = x$
5.	mul0	$x * 0 = 0$
6.	mul1	$x * 1 = x$
7.	subSelf	$x - x = 0$
8.	divSelf	$x / x = 1$
9.	subAdd	$x - y = x + (-y)$
10.	mulSym	$e * e' = e' * e$
11.	addSym	$e + e' = e' + e$
12.	mulAssoc	$(x * y) * z = x * (y * z)$
13.	addAssoc	$(x + y) + z = x + (y + z)$
14.	ifT	if True then lhs else rhs = lhs
15.	ifF	if False then lhs else rhs = rhs
16.	seqNop	seq lhs nop = lhs
17.	seqNop'	seq nop rhs = rhs
18.	repeatNop	repeat 0 body = nop
19.	repeatN1	repeat (S n) body = seq body (repeat n body)
20.	ifJoin	if c1 then (if c2 then op else nop) else nop = if (c1 and c2) then op else nop
21.	mapFusion	map g (map f xs) = map (g . f) xs
22.	foldMap	(foldr f e) . (map g) = foldr (f . g) e
23.	foldFusion	$\forall x y . h (f x y) = f' x (h y)$ $\implies h . (foldr f e) xs = foldr f' (h e) xs$

■ **Table 3** Basic Equality Axioms used in our Program Equivalence Case Study

Formula	Rewrites
$-(x + x) + (x + x) = 0$	7, 11, 9
$(x * 2) * 2 = (x + x + x + x)$	3, 13
$(x * y) + (y * x) = (x * 2 * y)$	3, 10, 12
$(x * y) + (y * z) - ((x + z) * y) = 0$	1, 7, 10
$(x * y) - (0 * y) = x * y$	2, 9, 7, 4
$x * (1 - (x/x)) = 0$	5, 7, 8
$x * 1 = x + 0$	4, 6
if true then (seq nop hw) else nop = hw	17, 14
repeat (S (S Z)) hw = seq hw hw	16, 18, 19
if True then (if False then hw else nop) else nop = if (True and False) then hw else nop	20
map p1 (map p2 list) = map p3 list	21
((foldr add 0) . (map p1)) list = foldr addP1 0 list	22
double . (foldr add 0) list = foldr twicePlus 0 list	23

■ **Table 4** Theorems Proved via Rewriting using the Basic Equality axioms in 3

D Dafny Matching Loop Example

```

datatype List = Nil | Cons(head: int, tail: List)

function append(xs: List, ys: List): List
{
  match xs
  case Nil => ys
  case Cons(x, xrest) => Cons(x, append(xrest, ys))
}

function reverse(xs: List): List
{
  match xs
  case Nil => Nil
  case Cons(x, xrest) => append(reverse(xrest), Cons(x, Nil))
}

lemma AppendNil(xs: List)
  ensures append(xs, Nil) == xs; {}

lemma AppendAssoc(xs: List, ys: List, zs: List)
  ensures append(xs, append(ys, zs)) == append(append(xs, ys), zs); {}

lemma ReverseAppendDistrib(xs: List, ys: List)
  ensures reverse(append(xs, ys)) == append(reverse(ys), reverse(xs));
{
  forall xs : List {AppendNil(xs);}
  forall xs : List, ys: List, zs: List {AppendAssoc(xs, ys, zs);}
}

lemma ReverseInvolution(xs: List)
  ensures reverse(reverse(xs)) == xs;
{
  // Axiom definition inserted here
  { forall (xs, ys) { ReverseAppendDistrib(xs, ys); } }

  match xs {
    case Nil =>
    case Cons(x, xrest) =>
      calc { // Equational reasoning steps removed here
        reverse(reverse(xs));
        {ReverseInvolution(xrest);}
        xs;
      }
  }
}

```

■ **Figure 11** A version of the reverse involution proof (Figure 9) from [34] with intermediate equality steps removed. Attempting to verify this code causes a matching loop when using Dafny version 3.3.0 and Z3 version 4.8.5