

Formulog = Datalog + ML + SMT

AARON BEMBENEK, Harvard University, USA

MICHAEL GREENBERG*, Pomona College, USA

STEPHEN CHONG, Harvard University, USA

Satisfiability modulo theories (SMT) solving has become a critical part of many static analyses, including symbolic execution, refinement type checking, and model checking. We propose Formulog, a domain-specific language that makes it possible to write a range of SMT-based static analyses in a way that is both close to their formal specifications and amenable to high-level optimizations and efficient evaluation.

Formulog extends the logic programming language Datalog with a first-order functional language and mechanisms for representing and reasoning about SMT formulas; a novel type system supports the construction of expressive formulas, while ensuring that neither normal evaluation nor SMT solving goes wrong. Our case studies demonstrate that a range of SMT-based analyses can naturally and concisely be encoded in Formulog, and that – thanks to this encoding – high-level Datalog-style optimizations can be automatically and advantageously applied to these analyses.

Additional Key Words and Phrases: Datalog, SMT, static analysis, refinement type systems, symbolic execution

1 INTRODUCTION

Satisfiability modulo theories (SMT) solving provides a way to reason logically about common program constructs such as arrays and bit vectors, and as such has become a key component of many static analyses. For example, symbolic execution tools use SMT solving to prune infeasible execution paths [Cadare et al. 2008; Cadare and Sen 2013]; type checkers use it to prove subtyping relations between refinement types [Bierman et al. 2012; Rondon et al. 2008]; and model checkers use it to abstract program states [Cimatti and Griggio 2012; McMillan 2006]. This paper presents Formulog, a domain-specific language for writing SMT-based static analyses. Formulog makes it possible to concisely encode a range of SMT-based static analyses in a way that is close to their formal specifications. Furthermore, Formulog is designed so that analyses implemented in it are amenable to efficient evaluation and powerful, high-level optimizations, including parallelization and automatic transformation of exhaustive analyses into goal-directed ones.

Formulog is based on Datalog, a logic programming language used to implement static analyses ranging from points-to analyses [Bravenboer and Smaragdakis 2009; Whaley and Lam 2004] to decompilers [Flores-Montoya and Schulte 2019; Grech et al. 2019] to security analyses [Grech et al. 2018; Guarnieri and Livshits 2009; Jordan et al. 2016; Livshits and Lam 2005; Tsankov et al. 2018]. Embodying the principle of separating the logic of a computation from the control necessary to perform that computation [Kowalski 1979], Datalog frees analysis designers from low-level implementation details and enables them to program at the level of specifications (such as formal inference rules). This leads to concise implementations [Whaley et al. 2005] that can be easier to reason about and improve at the algorithmic level compared to analyses in more traditional languages [Smaragdakis and Bravenboer 2011]. Datalog-based analyses can be fast and scalable, even outperforming the non-Datalog state-of-the-art [Bravenboer and Smaragdakis 2009]. Indeed, Datalog’s high-level

*Work done while on sabbatical at Harvard University.

Authors’ addresses: Aaron Bembenek, Harvard University, USA, bembenek@g.harvard.edu; Michael Greenberg, Pomona College, USA, michael@cs.pomona.edu; Stephen Chong, Harvard University, USA, chong@seas.harvard.edu.

2020. This is the author’s version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in .

nature makes it amenable to high-level optimizations, such as parallelization [Scholz et al. 2016] and synthesis of goal-directed analyses from exhaustive ones [Reps 1995].

However, despite the appeal of Datalog for static analysis and the importance of SMT solving in static analysis, until now there has not been a focused study of how to effectively extend the benefits of Datalog to SMT-based analyses; our work bridges this gap.

Formulog augments Datalog with an interface to an external SMT solver and a first-order fragment of the functional language ML. It provides a library of constructors for building terms that are interpreted as logical formulas when applied to special SMT operators; in the backend, these operators are implemented by calls to an external SMT solver. A Formulog program is essentially a set of ML-style function definitions and Datalog-style rules; both pieces can refer to each other and invoke the SMT operators. As in Datalog, the goal of Formulog evaluation is to compute all possible inferences with respect to the rules, which correspond to logical implications. Unlike Datalog, rule evaluation might involve both ML evaluation and calls to an SMT solver.

The way this design combines Datalog, ML, and SMT solving gives Formulog some desirable properties. First, Formulog programs can use SMT solving the way it is used in SMT-based analyses. This results from the choice to represent SMT formulas as ML terms, and contrasts with the approach of most prior work combining logic programming and constraint solving (where, e.g., checking for formula validity is hard). Second, the combination of Datalog-style rules and ML-style functions mirrors the combination of inference rules and helper functions commonly used in analysis specifications, making it easier to translate formal analysis specifications into executable code. This close correspondence between specification and implementation means that specification-level reasoning is still applicable to analysis implementations (and vice versa: unexpected behavior in Formulog programs has revealed bugs in specifications). Third, because Formulog is based on Datalog, analyses written in it can be effectively optimized and evaluated via powerful Datalog algorithms, making them competitive with analyses written in more mature languages.

It takes care to fit Datalog, ML, and SMT solving together in a way that truly achieves these properties. Along these lines, part of our technical contribution is a novel bimodal type system that treats terms appearing in SMT formulas more liberally than terms appearing outside of formulas, making it possible to construct expressive logical formulas, while still ensuring that neither concrete (i.e., Datalog/ML) evaluation nor SMT solving goes wrong.

To test the practicality of Formulog, we implemented a fully-featured, prototype Formulog runtime and wrote three substantial SMT-based analyses in Formulog: a type checker for a refinement type system, a bottom-up points-to analysis for JVM bytecode, and a bounded symbolic evaluator for a subset of LLVM bitcode. Our implementations for the first two case studies are almost direct translations of previously published formal specifications [Bierman et al. 2012; Feng et al. 2015]; indeed, Formulog allowed us to program close enough to the specifications to uncover bugs in both of them. Despite encoding complex analysis logic, each of our analyses is concise (no more than 1.5K LOC). Furthermore, our Formulog-based implementations have acceptable performance, even when compared against reference implementations running on more mature language platforms. In some cases, we actually achieve substantial speedups over the reference implementations.

These performance results are possible only because Formulog’s design allows our runtime to automatically and effectively apply high-level optimizations to Formulog programs. Our third case study makes this point emphatically. Due to automatic parallelization, our symbolic evaluator achieves a speedup of 7× over the symbolic execution tool KLEE [Cadar et al. 2008]. Moreover, this speedup increases to 12× when we use the magic set transformation [Bancilhon et al. 1985; Beeri and Ramakrishnan 1991] to automatically transform our exhaustive symbolic evaluator into a goal-directed one that explores only paths potentially leading to assertion failures. That Datalog

Programs	$\text{prog} ::= H^*$	Variables	$X \in \text{Var}$
Horn clauses	$H ::= p(e^*) :- P^*$	Constructors	$c \in \text{CtorVar}$
Premises	$P ::= A \mid !A$	Predicates	$p \in \text{PredVar}$
Atoms	$A ::= p(e^*) \mid e = e$		
Expressions	$e ::= X \mid c$		

Fig. 1. A Datalog program is a collection of Horn clauses that represent rules for making inferences.

can speed up analyses like points-to analysis is well established [Bravenboer and Smaragdakis 2009; Whaley and Lam 2004]; that it can automatically scale symbolic evaluation is a novel result.

In sum, this paper makes the following contributions:

- the design of Formulog (Section 3), a domain-specific language for writing SMT-based static analyses that judiciously combines Datalog, a fragment of ML, and SMT solving;
- a lightweight bimodal type system (Section 4) that mediates the interface between concrete evaluation and SMT solving, enabling the construction of expressive formulas while preventing many kinds of runtime errors in both concrete evaluation and SMT solving;
- a fully-featured prototype and three substantial case studies (Section 5), showing that the design of Formulog can be the basis of a practical tool for writing SMT-based analyses; and
- an evaluation of Formulog’s design in light of these case studies (Section 6), demonstrating how careful design decisions make Formulog an effective medium for encoding a range of SMT-based analyses in a way that is both close to their formal specifications and amenable to efficient evaluation and high-level optimizations.

2 BACKGROUND

The starting point for Formulog is Datalog with stratified negation (Figure 1) [Apt et al. 1988; Gallaire and Minker 1978; Green et al. 2013; Przymusiński 1988; Van Gelder 1989]. A Datalog program is a collection of Horn clauses, where a *clause* H consists of a head predicate $p(e^*)$ and a sequence of body premises P . Each *premise* P is either a positive atom A or a negated atom $!A$. An *atom* A has one of two forms: It is either a predicate symbol applied to a list of expressions, or the special equality predicate $e = e$. An *expression* e is a variable X or a nullary constructor c , i.e., an uninterpreted constant. Each predicate symbol p is associated with an extensional database (EDB) relation or an intensional database (IDB) relation. An *EDB relation* is tabulated explicitly through *facts* (clauses with empty bodies), whereas an *IDB relation* is computed through *rules* (clauses with non-empty bodies). A rule should be read as a universally quantified logical implication, with the conjunction of the body premises implying the predicate in the head. Datalog evaluation amounts to computing every possible inference with respect to these implications; the restriction to stratified negation (a relation cannot be defined, either directly or indirectly, by its complement) ensures that this can be done via a sequence of fixed point computations.

Datalog has proven to be a natural and effective way to encode a range of static analyses [Bravenboer and Smaragdakis 2009; Flores-Montoya and Schulte 2019; Grech et al. 2019, 2018; Guarnieri and Livshits 2009; Jordan et al. 2016; Livshits and Lam 2005; Tsankov et al. 2018; Whaley and Lam 2004]. EDB relations are used to represent the program under analysis; for example, EDB relations might encode a control flow graph of the input program. The logic of the analysis is encoded using rules that define IDB relations; these rules are fixed and do not depend on the program under analysis (which is already captured by the EDB relations). The Datalog program will compute the contents of the IDB relations, which can be thought of as the analysis results.

```

fun mem(X: 'a, Xs: 'a list) : bool =
  match Xs with [] => false | H :: T => X = H || mem(X, T) end

reaches(X, Y, [Y], Phi) :- (* reaches is an IDB relation *)
  edge(X, Y, Phi),          (* edge is an EDB relation, defined below *)
  is_sat(Phi) = true.

reaches(X, Z, Y :: Path, `Phi /\ Constraint`) :-
  edge(X, Y, Phi),
  reaches(Y, Z, Path, Constraint),
  mem(Y, Path) = false,
  is_sat(`Phi /\ Constraint`) = true.

edge(1, 2, `#x[bool]`). (* Constraint: SMT variable #x[bool] is true *)
edge(2, 3, `true`).    (* Constraint: true is true *)
edge(3, 4, `~#x[bool]`). (* Constraint: negation of #x[bool] is true *)

```

Fig. 2. This Formulog program computes a form of (non-reflexive) reachability for graphs where edges are labeled with propositions; here reachability requires the satisfiability of edge labels along a path.

That being said, standard Datalog is a very restricted language and there are many other analyses that cannot easily be encoded in it, if at all. Recent variants extend Datalog for analyses that operate over interesting lattices [Madsen et al. 2016; Szabó et al. 2018]. Following in this spirit, Formulog proposes a way to support analyses that need access to SMT solving.

3 LANGUAGE DESIGN

The design of Formulog is driven by three main desiderata. First, it should be possible to implement SMT-based static analyses in a form close to their formal specifications. Second, it should be easy to use logical terms the way that they are commonly used in many analyses. For example, analyses often need to create formulas about entities such as arrays and machine integers, test those formulas for satisfiability or validity, and generate models of them. Third, Formulog programs should still be amenable to powerful Datalog optimizations and evaluable using scalable Datalog algorithms.

Section 6 demonstrates how the design of Formulog largely meets these desiderata. Here, we give a warm-up example of Formulog, provide an overview of its language features, discuss how these features support logical formulas, and conclude with its operational semantics.

3.1 Formulog by example

To give the flavor of Formulog, we first present a small example (Figure 2). The “hello world” program of Datalog is graph transitive closure; here we give a Formulog analogue for directed graphs where edges are labeled with propositions, and we want to compute non-reflexive reachability modulo the satisfiability of the edge labels on a path.¹ The input to our program is an EDB relation *edge*, where $\text{edge}(x, y, \phi)$ denotes that there is an edge from node x to node y labeled with proposition ϕ . Our program computes an IDB relation *reaches*, where $\text{reaches}(x, y, p, \phi)$ denotes that node y is transitively reachable from node x through path $x : : p$ with cumulative constraint ϕ .

¹Computing reflexive reachability is also straightforward, but requires extra machinery (as it does for standard Datalog).

Types

Types	$\tau ::= t \mid t \text{ smt} \mid t \text{ sym} \mid \text{model}$
Pre-types	$t ::= B \mid D \tau^* \mid \alpha$
Base types	$B ::= \text{bool} \mid \text{string} \mid \text{bv}[k]_{k \in \mathbb{N}^+} \mid \dots$

Terms

Programs	$\text{prog} ::= H^* T^* F^* Z^*$
Horn clauses	$H ::= p(e^*) :- P^*$
Premises	$P ::= A \mid !A$
Atoms	$A ::= p(e^*) \mid e = e$
Type definitions	$T ::= \text{type } \alpha^* D = [c(\tau^*)]^*$
Functions	$F ::= \text{fun } f([X : \tau]^*) : \tau = e$
SMT declarations	$Z ::= \text{uninterpreted fun } c([t \text{ smt}]^*) : t \text{ smt} \mid$ $\text{uninterpreted sort } \alpha^* D$
Expressions	$e ::= X \mid c(e^*) \mid k \mid f(e^*) \mid \text{match } e \text{ with } [c(X^*) \rightarrow e]^* \mid$ $\text{let } X = e \text{ in } e \mid \text{if } e \text{ then } e \text{ else } e \mid \otimes(e^*) \mid \text{'}\phi\text{' } \mid p(w^*)$
Constants	$k ::= \text{true} \mid \text{false} \mid 0 \mid 1 \mid \dots$
SMT formulas	$\phi ::= e \mid c_{\text{forall}}^{\text{SMT}}(\phi, \phi) \mid c_{\text{let}}^{\text{SMT}}(\phi, \phi, \phi) \mid c_{\text{ctor}}^{\text{SMT}}[c](\phi^*) \mid \dots$
Wildcard	$w ::= ?? \mid e$
Values	$v \in \text{Val} ::= k \mid c(v^*)$

Namespaces

Data type names	$D \in \text{ADTVar}$	Variables	$X \in \text{Var}$
Type variables	$\alpha \in \text{TVar}$	Predicates	$p \in \text{PredVar}$
Constructors	$c \in \text{CtorVar}$	Functions	$f \in \text{FunVar}$

Fig. 3. Formulog extends the abstract syntax of Datalog with type definitions, functions, SMT declarations, and a richer language of expressions.

After using the ML fragment to define a simple function `mem` that decides list membership, our program states two rules defining our `reaches` relation. The first represents the base case: if there is an edge between nodes X and Y with a satisfiable label Phi , then Y is reachable from X along the path $[X, Y]$ with condition Phi . The built-in operator `is_sat` queries an external SMT solver for the satisfiability of its argument (an SMT proposition). The second rule computes the recursive case, making sure that the edge being used to extend the transitive closure is not redundant and that the labels along the prospective new path are satisfiable when taken together.

Given the edge relation in Figure 2, these rules compute that node 1 reaches nodes 2 and 3, node 2 reaches nodes 3 and 4, and node 3 reaches node 4, but node 1 does not reach node 4, since the path constraint `\#x[bool] /\ true /\ ~\#x[bool]` is not satisfiable.

3.2 Overview

Formulog extends Datalog with a fragment of first-order ML and a language of SMT formulas (Figure 3). Accordingly, a program consists of Horn clauses, type and function definitions, and SMT declarations. The Horn clause fragment is the same as in Datalog, except with a richer variety of expressions e that can occur as arguments to predicates.

Type definitions. Formulog users can define ML-style algebraic data types, which can be polymorphic and mutually recursive. An algebraic data type definition consists of a list of type variables α , a type name D , and a list of constructors c with their argument types τ . Section 4 explains Formulog’s type system in more detail; we provide a brief sketch now. Algebraic data types D , τ^* , base types B , and type variables α are treated as *pre-types*; intuitively, a pre-type t is the type of a concrete (non-formula) term. In addition to pre-types, there are types that represent SMT-relevant terms: a t -valued SMT formula has type t *smt*, a t -valued SMT variable has type t *sym*, and an SMT model — a finite map from formula variables to concrete terms — has type *model*. The Formulog type system distinguishes the first three types where it is computationally relevant (i.e., during concrete evaluation, where confusing a t -valued formula for a concrete t term might lead to a computation getting stuck), and collapses them where it is not (i.e., during SMT evaluation, where there is no meaningful distinction between a t -valued formula and a concrete t value). It also prevents SMT models, which are not representable as SMT expressions, from flowing into SMT formulas.

Functions. Formulog supports ML-style function definitions, although functions are limited to being first-order and are not first-class values. They can be polymorphic and mutually recursive.

SMT declarations. Formulog users can declare uninterpreted functions and polymorphic uninterpreted sorts. An uninterpreted function amounts to a special constructor for building a purely symbolic term of type t *smt* (for some pre-type t). An uninterpreted sort amounts to a special symbolic pre-type t , where t is not inhabited by any value, but t *sym* and t *smt* are.

Expressions and formulas. Expressions e occur as function bodies and as predicate arguments in Horn clauses. Although Datalog traditionally limits ground terms to nullary constructors, we admit n -ary constructors. While this comes with the cost of possibly-diverging programs — adding n -ary constructors makes Datalog Turing-complete [Green et al. 2013] — many recent Datalog variants allow complex terms, including Soufflé [Scholz et al. 2016], LogicBlox [Aref et al. 2015], and Flix [Madsen et al. 2016]. For us, complex terms provide a natural way to reify logical formulas, and they also can be used to create data structures that make it easier to encode certain analyses.

Additional Formulog expressions include standard ML fare like constants (booleans, strings, machine integers, and floats), function calls, and *match*, *let*, and *if-then-else* expressions. The expression $\otimes(e^*)$ represents the application of a primitive operator to a sequence of subexpressions. These cover both basic arithmetic operations (e.g., addition) and SMT-specific operations (e.g., checking for satisfiability, generating models; see Section 3.3.2).

The expression $\` \phi \`$ is a quasi-quoted SMT formula, where the language of formulas ϕ consists of unquoted expressions e and formula constructors of the form c_e^{SMT} applied to SMT formulas. Some of these constructors directly reflect SMT formula constructs; for example, the constructor $c_{\text{forall}}^{\text{SMT}}$ builds a universally quantified formula, and the constructor $c_{\text{let}}^{\text{SMT}}$ builds an SMT *let* formula. Formula constructors can appear only in formulas, and non-formula constructors cannot appear directly in formulas. We embed algebraic data type constructors in formulas using a family of *formula constructors*. Each formula constructor $c_{\text{ctor}}^{\text{SMT}}[c]$ lifts the user-defined algebraic data type constructor c to SMT. Quotes are used to delineate formulas and trigger a different type checking mode, in which the types t , t *smt*, and t *sym* are conflated (with some restrictions, as explained in Section 4). The unquote operator $,$ escapes from this type checking mode and makes it possible to inject a non-formula expression into a formula. Section 3.3 discusses formulas in more detail.

We have already seen how the Datalog fragment of Formulog can include expressions from the ML fragment; the final expression $p(w^*)$ ties the loop by providing a way for the ML fragment to reference the Datalog fragment. The expression $p(w^*)$ acts like a function call that queries the contents of the relation p . Its exact behavior depends on its arguments, which are either

Negation	\sim	: $\text{bool smt} \rightarrow \text{bool smt}$
Conjunction	\wedge	: $(\text{bool smt}, \text{bool smt}) \rightarrow \text{bool smt}$
Implication	\implies	: $(\text{bool smt}, \text{bool smt}) \rightarrow \text{bool smt}$
Equality	$\#=[t]$: $(t \text{ smt}, t \text{ smt}) \rightarrow \text{bool smt}$
SMT variable	$\#\{\cdot\}[t]$: $'a \rightarrow t \text{ sym}$
Bit vector constant	$\text{bv_const}[k]$: $\text{bv}[32] \rightarrow \text{bv}[k] \text{ smt}$
Bit vector addition	bv_add	: $(\text{bv}[k] \text{ smt}, \text{bv}[k] \text{ smt}) \rightarrow \text{bv}[k] \text{ smt}$

Fig. 4. Logical formulas are created in Formulog via built-in constructors, such as the ones shown here.

expressions or the special wildcard term `??`. If its arguments contain no wildcards, then $p(e^*)$ returns a boolean indicating whether the tuple identified by its arguments is in the p relation. If it has $k > 0$ wildcards, it returns a list of k -tuples: For each tuple v^* in the relation corresponding to p , there is a corresponding k -tuple in this list that is v^* projected to the wildcard positions; if there are n matching tuples in p , then the list is of length n . In other words, given complete arguments, a predicate is really just a predicate; given partial arguments with wildcards, a predicate is the multiset consisting of matching tuples after they have been appropriately projected.

Remarks. Extending Datalog with our fragment of ML is not foundational, as it can relatively easily be translated to Datalog rules (this would not necessarily be the case for a higher-order fragment of ML). However, despite the fact that the ML fragment could be treated as just syntactic sugar, it has a significant positive impact on the usability of Formulog, as we argue in Section 6.

The concrete syntax of formulas in our prototype (and in the examples we give in this paper) differs from the abstract syntax given here. First, algebraic data type constructors are allowed to appear directly in formulas, and are implicitly lifted to the appropriate formula constructor (so data type constructor c is automatically lifted to $c_{\text{ctor}}^{\text{SMT}}[c]$). Second, we do not support an explicit unquote operator; instead, we implicitly unquote variables, constants, and invocations of nullary functions.

3.3 Logical formulas

Formulog uses data types and operators to support constructing and reasoning about logical formulas. Formulog provides a library of data types that define logical terms. Most of the time during evaluation, these terms are unremarkable and treated just like any other ground term. However, these terms are interpreted as logical formulas when they are used as arguments to built-in operators that make calls to an external SMT solver. In our current prototype, it is possible to create logical terms in first-order logic extended with (fragments of) the SMT-LIB theories of uninterpreted functions, integers, bit vectors, floating point numbers, arrays, and algebraic data types [Barrett et al. 2017], as well as the theory of strings shared by the SMT solvers Z3 [de Moura and Bjørner 2008] and CVC4 [Barrett et al. 2011].

3.3.1 Representing formulas. Users create logical terms through constants and formula constructors. For example, to represent the formula $\text{False} \implies \text{True}$, one would use the term ``false ==> true``, where `false` and `true` are the standard boolean values and `==>` is the infix constructor for implication.

Formulog offers around 70 constructors for creating logical terms ranging from symbolic string concatenation to logical quantifiers. Figure 4 shows a sample of these constructors and their types. Some constructors require explicit indices, either to guarantee that type information is available at runtime when the formula is serialized to SMT-LIB, or to make sure that the type of the arguments can be determined by the type of the constructed term (which makes type checking easier). For example, `bv_const[k]` creates a symbolic k -bit vector value from a concrete 32-bit vector; at runtime,

Satisfiability	<code>is_sat</code>	: <code>bool smt</code> \rightarrow <code>bool</code>
	<code>is_sat_opt</code>	: <code>(bool smt list, bv[32] option)</code> \rightarrow <code>bool option</code>
Validity	<code>is_valid</code>	: <code>bool smt</code> \rightarrow <code>bool</code>
Model generation	<code>get_model</code>	: <code>(bool smt, bv[32] option)</code> \rightarrow <code>model option</code>
Model inspection	<code>query_model</code>	: <code>('a sym, model)</code> \rightarrow <code>'a option</code>

Fig. 5. Formulog provides built-in operators for reasoning about logical terms.

it is necessary to know the width k so that we can serialize it correctly. On the other hand, in the case of the constructor `#=[t]`, denoting the equality of two terms of type t `smt`, the index makes sure that the type checker knows what types the arguments should have. A programmer typically does not need to provide these indices explicitly, as they can often be inferred (our prototype does this).

Formulog distinguishes between logic programming variables and formula variables. A formula variable is a ground term that, when interpreted logically, represents a symbolic value. The term `#{v}[t]` is a formula variable of type t `sym` identified by the value v (which can be of arbitrary type). Intuitively, v is the “name” of the variable. The term `#{v}[t]` is guaranteed not to occur in v , which means that the variable it represents is fresh with respect to the set of formula variables in v ; this makes it easy to deterministically construct a new variable that is fresh with respect to an environment, a trick we use often in our case studies. For example, if χ is bound to a list of boolean formula variables, the formula variable `#{X}[bool]` will not unify with any term in χ . The shorthand `#id[t]` is equivalent to `#{"id"}[t]`, where `id` is a syntactically valid identifier.

Importantly, because formula variables are ground terms, we can derive facts containing formula variables without violating Datalog’s range restriction, which requires that every derived fact is variable-free. This restriction enables efficient evaluation by simplifying table lookups, one of the fundamental operations in Datalog evaluation.

3.3.2 Using formulas. Built-in operators provide a way to reason about logical terms as formulas (Figure 5). When an operator in the SMT interface is invoked, its formula argument is serialized into the SMT-LIB format and a call is made to an external SMT solver. These operators are assumed to act deterministically during a single Formulog run; an implementation can achieve this in the presence of a non-deterministic SMT solver by memoizing operations.

For example, to test the validity of the principle of explosion (any proposition follows from false premises), one could make the call `is_valid(`false ==> #x[bool]`)`. Like other operators, the SMT interface operators can be invoked from the bodies of rules, as here:

```
ok :- #x[bool] != #y[bool],
      is_sat(`#x[bool] #= #y[bool]`) = true,
      is_sat(`~(#x[bool] #= #y[bool])`) = true.
```

This rule derives the fact `ok`: The term `#x[bool]` is not unifiable with the term `#y[bool]`, since they are different formulas, representing different SMT variables. But these terms both may and may not be equal when interpreted as formula variables via the operator `is_sat`. Within an invocation of `is_sat`, constraints are formed between `#x[bool]` and `#y[bool]` — in the first case they must be equal, and in the second case they must not be — but these constraints do not leak into the larger context. This is an intentional design decision and differs from the approach taken by paradigms like constraint logic programming (see Section 6).

Formulog provides two sets of operators for testing the satisfiability and logical validity of propositions. In general, an SMT solver can return three possible answers to such a query: “yes,” “no,” and “unknown.” The operators `is_sat` and `is_valid` return booleans. In the case that the

backend SMT solver is not be able to determine whether a formula ϕ is satisfiable, these operators fail (as explained in Section 4). The operator `is_sat_opt(ϕ^* , timeout)` provides more fine-grained control: it takes a list of propositions (interpreted as conjuncts) and an optional timeout, and returns an optional boolean, with `none` corresponding to “unknown.” While we suspect that the simpler versions will be sufficient for most applications, this more complex version does allow applications to explicitly handle the “unknown” case if need be (e.g., pruning paths in symbolic execution).

The operator `get_model` takes a proposition and an optional timeout; it returns a model for the proposition if the SMT solver is able to find one in time, and `none` otherwise. The values of formula variables in this model can be inspected using `query_model`, which returns `none` if the variable does not occur free in the formula or if a concrete value for it is not representable in Formulog (for example, Formulog does not have a type for a concrete 13-bit vector). The values of symbolic expressions can be indirectly extracted through formula variables: Before finding the model, add the equality ``x #= e`` to the formula, where x is a fresh formula variable and e is an expression; in the extracted satisfying model, x will be assigned the value of e in that model.

3.3.3 Custom types in formulas. Formulog’s algebraic data types can be reflected in SMT formulas via SMT-LIB’s support for algebraic data types. Thus, Formulog permits arbitrary term constructors to be used within logical formulas. For example, we can define a type `foo` with a single nullary constructor `bar` and then write formulas involving `foo`-valued terms:

```
type foo = | bar
ok :- is_valid(`#x[foo] #= bar`) = true.
```

This program would derive the fact `ok`: Since there is only one way to construct a `foo` — through the constructor `bar` — any symbolic value of type `foo` must be the term `bar`.

For each algebraic data type, we automatically generate two kinds of constructors that make it easier to write formulas involving terms of that type. The first kind is a constructor tester. For each constructor c of a type t , Formulog provides a constructor `#is_c` of type $t \text{ smt} \rightarrow \text{bool smt}$. The proposition `#is_c(e)` holds if the outermost constructor of e is c . The second kind is an argument getter. If c is a constructor for type t with n arguments of types t_i for $1 \leq i \leq n$, Formulog generates n argument getters of the form `#c_i`, where `#c_i` has the type $t \text{ smt} \rightarrow t_i \text{ smt}$. When interpreted as a formula, the term `#c_i(e)` represents the value of the i^{th} argument of e . For example, we can state that a symbolic list of booleans is non-empty and its first argument is `true`:

$$\text{\`#is_cons(\#x[bool list])} \wedge \text{\#cons_1(\#x[bool list])\`}$$

We could use the operator `get_model` to find a model of this satisfiable formula; in this model, `#x[bool list]` might be assigned the concrete value `cons(true, nil)`.

3.4 Operational semantics

This section presents Formulog’s operational semantics, making reference to a selection of the formal rules (Figures 6).^{2,3} Formulog imposes the standard stratification requirements upon programs: no recursive dependencies involving negation or aggregation between relations. As a stratifiable program can be evaluated one stratum at a time, we focus on the evaluation of a single stratum.

A stratum is evaluated by repeatedly evaluating its Horn clauses until no new inferences can be made. The semantics of a Horn clause H is defined through the judgment $\vec{F}; \mathcal{W} \vdash H \rightarrow \mathcal{W}_L$, where

²Formulog can also be given a model-theoretic semantics: because the ML features can be desugared into Datalog rules, the model theory of Formulog is essentially that of stratified Datalog. Appendix F sketches this out further.

³In the boxed rule schemata, implicit parameters are in `gray`; we conserve space by stating the rules without threading implicit parameters through, which are unchanging. We write \vec{x}_i for some metavariable x to mean a possibly empty sequence of x s indexed by i , and write S_\perp for some set S to mean the set $S + \text{Err}$.

Namespaces and constructs

World $\mathcal{W} \in \text{PredVar} \rightarrow \mathcal{P}(\text{Val}^*)$
 Substitution $\theta \in \text{Var} \rightarrow \text{Val}$

Error $\perp \in \text{Err}$
 u -term $u ::= X \mid k \mid c(\vec{u}_i)$

Clause semantics

$$\vec{F}; \mathcal{W} \vdash H \rightarrow \mathcal{W}_\perp$$

$$\frac{|\vec{P}_i| = n \quad \theta_0 = \cdot \quad \forall i \in [0, n), \theta_i \vdash P_i \rightarrow \theta_{i+1}}{\vec{F}; \mathcal{W} \vdash p(\vec{X}_j) :- \vec{P}_i \rightarrow \mathcal{W}[p \mapsto \mathcal{W}(p) \cup \{\theta_n(\vec{X}_j)\}]} \text{CLAUSE}$$

Premise semantics

$$\vec{F}; \mathcal{W}; \theta \vdash P \rightarrow \theta_\perp$$

$$\frac{\vec{v} \in \mathcal{W}(p) \quad \theta \vdash \vec{X} \sim \vec{v} : \theta'_\perp}{\mathcal{W}; \theta \vdash p(\vec{X}) \rightarrow \theta'_\perp} \text{PosATOM}$$

$$\frac{\theta \vdash Y \sim c(\vec{X}) : \theta'_\perp}{\mathcal{W}; \theta \vdash Y = c(\vec{X}) \rightarrow \theta'_\perp} \text{EqCTOR}$$

Expression semantics

$$\vec{F}; \mathcal{W}; \theta \vdash e \Downarrow_e v_\perp$$

$$\vec{F}; \mathcal{W}; \theta \vdash \vec{e} \Downarrow_{\vec{e}} \vec{v}_\perp$$

$$\frac{\mathcal{W}; \theta \vdash \vec{e} \Downarrow_{\vec{e}} \vec{v} \quad \llbracket \otimes \rrbracket(\vec{v}) = v}{\mathcal{W}; \theta \vdash \otimes(\vec{e}) \Downarrow_e v} \Downarrow_e\text{-OP}$$

$$\frac{\mathcal{W}; \theta \vdash \phi \Downarrow_\phi v_\perp}{\mathcal{W}; \theta \vdash \text{'}\phi\text{'} \Downarrow_e v_\perp} \Downarrow_e\text{-QUOTE}$$

Formula semantics

$$\vec{F}; \mathcal{W}; \theta \vdash \phi \Downarrow_\phi v_\perp$$

$$\vec{F}; \mathcal{W}; \theta \vdash \vec{\phi} \Downarrow_{\vec{\phi}} \vec{v}_\perp$$

$$\frac{\mathcal{W}; \theta \vdash \vec{\phi} \Downarrow_{\vec{\phi}} \vec{v}}{\mathcal{W}; \theta \vdash c_c^{\text{SMT}}(\vec{\phi}) \Downarrow_\phi c_c^{\text{SMT}}(\vec{v})} \Downarrow_\phi\text{-CTOR}$$

$$\frac{\mathcal{W}; \theta \vdash e \Downarrow_e v}{\mathcal{W}; \theta \vdash ,e \Downarrow_\phi \text{toSMT}(v)} \Downarrow_\phi\text{-UNQUOTE}$$

SMT conversion

$$\text{toSMT}(v) = v$$

$$\begin{aligned} \text{toSMT}(c_{\text{let}}^{\text{SMT}}(v_1, v_2, v_3)) &= c_{\text{let}}^{\text{SMT}}(v_1, v_2, v_3) \\ \text{toSMT}(c_{\text{forall}}^{\text{SMT}}(v_1, v_2)) &= c_{\text{forall}}^{\text{SMT}}(v_1, v_2) \end{aligned}$$

$$\begin{aligned} \text{toSMT}(c(\vec{v}_i)) &= c_{\text{ctor}}^{\text{SMT}}[c](\overrightarrow{\text{toSMT}(v_i)}) \\ &\dots \end{aligned}$$

Fig. 6. A fragment of Formulog's operational semantics (see Appendix C for full formalization).

a world \mathcal{W} is a map from predicate symbols to sets of tuples (i.e., those that have been derived so far). A Horn clause takes a world to either a new world or the error value \perp . Going wrong can result for two reasons: either because a variable is unbound at a point where it needs to be bound, or because an operator is applied to a value outside of its domain. It is important to distinguish between a rule going wrong and a rule failing to complete because two terms fail to unify: The first is an undesirable error (ruled out by our type system), whereas the second is expected behavior.

A rule is evaluated by evaluating its premises one-by-one, using a left-to-right order (CLAUSE). The judgment $\vec{F}; \mathcal{W}; \theta \vdash P \rightarrow \theta_\perp$ defines the semantics of a premise, which takes a world and a substitution θ (a partial function from variables to values) and returns a new substitution or an error. The substitution produced by one premise is used as the input to the next one. A successful inference extends the input world with a (potentially novel) tuple $\theta_n(\vec{X}_j)$, i.e., the result of element-wise applying the substitution produced by the rightmost premise to the variables in the head of the rule. Clause evaluation goes wrong if the evaluation of one of the premises goes wrong.

Without loss of generality, we assume that premises occur in a limited form: predicates are applied to only variables, written $p(\vec{X}_i)$, and equality predicates bind variables, as in $Y = e$. (Our prototype similarly desugars premises.) An atom $p(\vec{X})$ is evaluated by non-deterministically choosing a tuple \vec{v} from the tuples in $\mathcal{W}(p)$, and then pairwise unifying its elements with the variables \vec{X} (PosAtom). The premise $Y = c(\vec{X})$ unifies its two terms (EQCTOR). The judgment $\theta \vdash u_1 \sim u_2 : \theta_\perp$ defines the unification of terms u_1 and u_2 under the substitution θ ; it results in an error if u_1 and u_2 both contain unbound variables, and a new substitution if they are otherwise unifiable.

The semantics for many Formulog expressions are standard. The evaluation of an operator produces a value if its arguments are evaluated to values in its domain (\Downarrow_e -OP); it goes wrong if the argument values are outside of its domain, e.g., if a string and number are added together. A quoted formula $\texttt{'}\phi\texttt{'}$ evaluates to whatever ϕ evaluates to (\Downarrow_e -QUOTE). Formula $c_c^{\text{SMT}}(\vec{\phi})$ evaluates to formula $c_c^{\text{SMT}}(\vec{v})$ if arguments $\vec{\phi}$ evaluate to values \vec{v} (\Downarrow_ϕ -CTOR). If the expression e evaluates to the value v , then the formula $\texttt{,}e\texttt{,}$ evaluates to the term $\text{toSMT}(v)$ (\Downarrow_ϕ -UNQUOTE), where the helper function toSMT lifts a term to its formula version.

4 TYPE SYSTEM

Formulog's type system is designed to meet three desiderata. The first desideratum is that concrete evaluation should never go wrong, which might happen if an operator is applied to an operand outside its domain or a variable is unbound at a point when it needs to be evaluated. The second desideratum is that SMT solving should never go wrong, which might happen if a term that does not represent a well-sorted formula under the SMT-LIB standard reaches the external SMT solver (e.g., a formula representing the addition of a 16-bit vector and 32-bit vector). The third desideratum is that the type system should make it easy to construct expressive logical formulas, including formulas that involve terms drawn from user-defined types.

There is some tension between the first and third of these desiderata. The first one requires that we differentiate between, for example, a concrete bit vector value and a symbolic bit vector value (e.g., a bit vector-valued formula) since an operator that is expecting a concrete bit vector might get stuck if its argument is a symbolic bit vector. For instance, we want to rule out this program:

Example 1 (A bad program we would want to reject).

```
type foo = | bar(bv[32])
fun f(F: foo) : bv[32] = match F with bar(Y) => Y + Y end
not_ok :- X = #x[bv[32]],
          f(bar(X)) = 42.
```

This program gets stuck evaluating $f(\text{bar}(X))$, since Y is bound to a symbolic value in f but the ML fragment's addition operator needs concrete arguments. On the other hand, we are able to construct more expressive formulas if we can occasionally conflate concrete and symbolic expressions:

Example 2 (A good program we would want to accept).

```
ok :- X = #x[bv[32]],
      is_sat(`bar(X) #= bar(5)`) = true.
```

This rule asks whether there exists a symbolic bit vector x such that $\text{bar}(x)$ equals $\text{bar}(5)$, where bar is the constructor defined above. This reasonable formula is not well-typed under a type system that uniformly distinguishes between concrete and symbolic values, since the constructor bar expects a concrete bit vector argument but instead receives the symbolic one x .

Contexts

Data type declarations	$\Delta ::= \cdot \mid \Delta, D : \forall \vec{\alpha}_i. \overrightarrow{\{c_j : \vec{\tau}_k\}}$
Program declarations	$\Phi ::= \cdot \mid \Phi, f : \forall \vec{\alpha}, \vec{\tau} \rightarrow \tau \mid \Phi, p \subseteq \vec{\tau}$
Variable contexts	$\Gamma ::= \cdot \mid \Gamma, x : \tau \mid \Gamma, \alpha$

Clause typing

$$\boxed{\Delta; \Phi \vdash H}$$

$$\frac{\cdot \vdash P_0 \triangleright \Gamma_1 \quad \dots \quad \Gamma_j \vdash P_j \triangleright \Gamma_{j+1} \quad \dots \quad \Gamma_n \vdash P_n \triangleright \Gamma' \quad p \subseteq \vec{\tau}_i \in \Phi \quad \Gamma' \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'}{\vdash p(\vec{X}_i) :- \vec{P}_j} \quad H\text{-CLAUSE}$$

Variable binding and typing

$$\boxed{\Gamma \vdash x, \tau \triangleright \Gamma}$$

$$\boxed{\Gamma \vdash \vec{x}, \vec{\tau} \triangleright \Gamma}$$

$$\frac{X \notin \text{dom}(\Gamma)}{\Gamma \vdash X, \tau \triangleright \Gamma, X : \tau} \quad X\tau\text{-BIND}$$

$$\frac{\Gamma(X) = \tau}{\Gamma \vdash X, \tau \triangleright \Gamma} \quad X\tau\text{-CHECK}$$

Premise typing

$$\boxed{\Delta; \Phi; \Gamma \vdash P \triangleright \Gamma}$$

$$\frac{p \subseteq \vec{\tau}_i \in \Phi \quad \Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'}{\Gamma \vdash p(\vec{X}_i) \triangleright \Gamma'} \quad P\text{-PosATOM}$$

$$\frac{\Gamma \vdash e : \tau \quad \Gamma \vdash Y, \tau \triangleright \Gamma'}{\Gamma \vdash Y = e \triangleright \Gamma'} \quad P\text{-EQ-FB}$$

Function and expression well formedness

$$\boxed{\Delta; \Phi \vdash F}$$

$$\boxed{\Delta; \Phi; \Gamma \vdash e : \tau}$$

$$\frac{\text{typeof}(\otimes) = \vec{\tau}_i \rightarrow \tau \quad \Gamma \vdash e_i : \tau_i}{\Gamma \vdash \otimes(\vec{e}_i) : \tau} \quad e\text{-OP}$$

$$\frac{\Gamma \vdash \phi : \tau}{\Gamma \vdash \text{'}\phi\text{'}} \quad e\text{-QUOTE}$$

SMT constructors and formula well formedness

$$\boxed{\Delta; \Phi; \Gamma \vdash c_{\dots}^{\text{SMT}} : \vec{\tau}_i \rightarrow \tau}$$

$$\boxed{\Delta; \Phi; \Gamma \vdash \phi : \tau}$$

$$\frac{\Gamma \vdash c_c^{\text{SMT}} : \vec{\tau}_i \rightarrow \tau \quad \Gamma \vdash \phi_i : \tau_i}{\Gamma \vdash c_c^{\text{SMT}}(\vec{\phi}_i) : \tau} \quad \phi\text{-CTOR}$$

$$\frac{\Gamma \vdash e : \tau}{\Gamma \vdash ,e : \text{toSMT}(\tau)} \quad \phi\text{-UNQUOTE}$$

$$\frac{\Gamma \vdash \phi : t \text{ sym}}{\Gamma \vdash \phi : t \text{ smt}} \quad \phi\text{-PROMOTE}$$

SMT representations

$$\boxed{\text{erase}(\tau) = t}$$

$$\boxed{\text{toSMT}(\tau) = \tau}$$

$$\begin{array}{lll} \text{erase}(B) = B & \text{erase}(t \text{ smt}) = \text{erase}(t) & \text{toSMT}(t) = \text{erase}(t) \text{ smt} \\ \text{erase}(D \vec{\tau}_i) = D \overrightarrow{\text{erase}(\vec{\tau}_i)} & \text{erase}(t \text{ sym}) = \text{erase}(t) & \text{toSMT}(t \text{ smt}) = \text{erase}(t) \text{ smt} \\ & & \text{toSMT}(t \text{ sym}) = \text{erase}(t) \text{ sym} \end{array}$$

Fig. 7. A fragment of Formulog's type system (see Appendix B for full formalization).

Formulog resolves the tension between these desiderata through a bimodal type system that acts differently inside and outside formulas (which are demarcated by quotations). In essence, the Formulog type system differentiates between the pre-type t , the SMT formula type $t \text{ smt}$, and the SMT variable type $t \text{ sym}$ outside of formulas, but typically conflates them within formulas.⁴

⁴It does not conflate them in binding positions where formula variables are required, such as in quantifiers.

This bimodal approach disallows Example 1 (since outside a formula, a term of type $\text{bv}[32]$ sym cannot be used where a term of type $\text{bv}[32]$ is expected), while permitting Example 2 (since within a formula, a term of type $\text{bv}[32]$ sym can be used anywhere a term of type $\text{bv}[32]$ is expected).

Intuitively, this bimodal approach is safe because it distinguishes between concrete and symbolic values during concrete evaluation – where conflating them might lead to going wrong – and conflates them only during SMT evaluation, where the distinction is not meaningful. We have formalized the Formulog type system and proven it sound with respect to the operational semantics of Formulog. We present only a small subset of it here (Figure 7); the full system is in Appendix B.

The rule defining a well-typed Horn clause (H -CLAUSE) depends on two notable judgments. The premise typing judgment $\Gamma \vdash P \triangleright \Gamma'$ takes a variable typing context Γ and a premise P and produces a new variable typing context Γ' . The variable binding and typing judgment $\Gamma \vdash x, \tau \triangleright \Gamma'$ holds if either X is not in Γ , in which case Γ' extends Γ with X mapped to τ ($X\tau$ -BIND), or X is mapped to τ by Γ , in which case $\Gamma = \Gamma'$ ($X\tau$ -CHECK). As can be seen from rule H -CLAUSE, the type checking of Horn clauses is flow-sensitive and proceeds left-to-right across the clause, with the “output” context of checking premise P_i used as the “input” context for checking premise P_{i+1} . This left-to-right type checking mirrors the left-to-right evaluation strategy Formulog uses; this is important for ensuring that variables are bound at the correct points.⁵ The second line of premises in rule H -CLAUSE ensures that every variable in the head of the rule is bound at the type specified by the head relation’s signature.

A positive atom is well typed if each of its variable arguments has the type given to that argument by the relation’s signature (P -PosATOM). A premise of the form $Y = e$ is typed according to a few different rules depending on which side of the equation is ground with respect to the input context Γ . The key is that our type system only types premises of the form $Y = e$ when unification is guaranteed to not go wrong at runtime.

The typing rules for most expressions are standard. An operation is well-typed if its arguments match its type signature (ϕ -CTOR). A quoted formula ϕ types at whatever ϕ types at (e -QUOTE). The formula constructor c_c^{SMT} is well typed if the types of its arguments match its type signature (ϕ -CTOR); in the case of a constructor for an algebraic data type that has been lifted to a formula constructor, that signature will require the constructed term and all of its arguments to have types of the form t smt. If an expression types at τ , then the formula e types at $\text{toSMT}(\tau)$ (ϕ -UNQUOTE). The helper function toSMT lifts a type to a formula type; for example, it lifts bool to bool smt . The typing rules for formulas also include a rule promoting from t sym to t smt, reflecting the fact that, within a formula, a t -valued formula variable can be used anywhere a t -valued formula can be.⁶

Type soundness with respect to the semantics of Formulog comes from safety and preservation:

THEOREM 4.1 (SAFETY). *If $\Delta; \Phi \vdash \vec{F}_i \vec{H}_j$ and $\Delta; \Phi \models \mathcal{W}$ then for all $H \in \vec{H}_j$, $\neg(\vec{F}_i; \mathcal{W} \vdash H \rightarrow \perp)$.*

THEOREM 4.2 (PRESERVATION). *If $\Delta; \Phi \vdash \vec{F}_i \vec{H}_j$ and $\Delta; \Phi \models \mathcal{W}$ and $\vec{F}_i; \mathcal{W} \vdash H \rightarrow \mathcal{W}'$ for some $H \in \vec{H}_j$ then $\Delta; \Phi \models \mathcal{W}'$.*

Safety (Theorem 4.1) guarantees that a Horn clause from a well-typed program, evaluated on a well-typed world (i.e., one where all the tuples have the right types), cannot step to error. Thus, safety means that an operator is never applied to an operand outside its domain, and a variable is never unbound when it needs to be bound. Preservation (Theorem 4.2) guarantees that if a Horn

⁵The fact that the operational semantics and type system assume a certain order of evaluation does not prohibit a Formulog runtime from reordering premises within rules (for example, when applying database-style query planning optimizations); it just needs to check that the new order is also well typed. This type of rewriting does not affect the result of running the rule provided that all subexpressions terminate (an assumption we make).

⁶The opposite is not true, since some formula constructors (i.e., quantifiers and let expressions) bind formula variables.

clause, from a well-typed program, is evaluated on a well-typed world and results in a new world, then that new world is also well-typed. Taken together, these theorems imply that a well-typed Formulog program does not go wrong during concrete evaluation (see Appendix E for proofs).

The type system is sound with respect to the semantics of SMT-LIB because the types of the formula constructors provided by Formulog are consistent with the types given by the SMT-LIB standard. The Formulog type system guarantees that, at runtime, terms (including formulas) are well-typed, and the type system prevents terms that are not representable in SMT (such as those of type model) from flowing into SMT formulas. We distinguish between SMT-compatible types and non-SMT types formally by indexing the type well formedness judgment with a *mode*, which is either *smt* (for those types that can be sent to the solver) or *exp* (for those types that cannot). It is fair to think of these modes as kinds with a subkinding relationship: types of kind *smt* can safely be treated as general types of kind *exp*, but not the other way round (Lemma D.1).

Assumptions. An actual implementation of Formulog, such as our prototype, has to contend with a few sources of going wrong that are not captured in our formal model. First, our model assumes that patterns in match clauses are exhaustive; this is just for simplicity, and could be statically checked using standard algorithms. Second, our model assumes that operators are total with respect to terms with the correct type. There are three places where this assumption might break: 1) division or remainder by zero; 2) the operators `is_sat` or `is_valid` may induce an “unknown” response from the external SMT solver; and 3) the SMT solver may reject patterns used in trigger-based quantifier instantiation that it considers to be ill formed (for example, if the pattern contains a binding operation). The first case is standard for many languages; the second can be avoided if the programmer uses the option-returning SMT operator `is_sat_opt`. The last case would be hard to check statically; however, an implementation could dynamically check patterns before making a call to the SMT solver, dropping invalid patterns and issuing a warning to the user. Our prototype uses “hard exceptions” by default, aborting the program. We also support a “soft exception” mode, which treats all these cases analogously to unification failures, halting execution on the current path but allowing execution on other paths to continue.

5 IMPLEMENTATION AND CASE STUDIES

In this section we briefly describe our prototype implementation of Formulog, and then discuss three analyses we have built as case studies: refinement type checking, bottom-up points-to analysis, and bounded symbolic evaluation.

5.1 Prototype

Our prototype runtime (~17.5K LOC Java) works in five stages: parsing, type checking, rewriting (for query specialization), validation, and evaluation. Stratification and the range restriction are checked during the validation phase. Our parallel implementation of semi-naive evaluation [Bancilhon 1986] uses a work-stealing thread pool; worker threads dispatch SMT queries to a pool of Z3 instances [de Moura and Bjørner 2008]. Our prototype is feature complete, but not heavily optimized.

Unless otherwise noted, we ran experiments on an Ubuntu Server 16.04 LTS machine with a 3.1 GHz Intel Xeon Platinum 8175 processor (24 physical CPUs, each hyperthreaded) and 192 GiB of memory. We configured our Formulog runtime to use up to 40 threads and up to 40 Z3 instances (v4.8.7); all comparison systems were set to use the same version of Z3 (with one exception, noted later). For each result, we report the median of three trials.⁷ Times are given as minutes:seconds.

⁷For each case study, we use a tool to take a program in the input language (e.g., Java) and turn it into Formulog facts. We do not include these times, as they are typically quite short. Extracting facts from libraries can take a few minutes, but this needs to be done only once per library.

```

fun accum_nil_axiom : bool smt =
  let (F, I) = (#func[closure], #init[enc_val]) in
  `forall F, I : accum(F, v_zero, I). accum(F, v_zero, I) #= I`

```

Fig. 8. This axiom encodes the denotation of a Dminor accumulate expression over an empty multiset. The term in the formula between `:` and `.` is a quantifier pattern [Detlefs et al. 2005].

5.2 Refinement type checking

We have implemented a type checker in Formulog for Dminor, a first-order functional programming language for data processing that combines refinement types with dynamic type tests [Bierman et al. 2012]. This type system can, e.g., prove that

$$x \text{ in } \text{Int} \ ? \ x : (x \ ? \ 1 : \emptyset)$$

type checks as `Int` in a context in which `x` has the union type $(\text{Int} \mid \text{Bool})$. Proving this entails encoding types and expressions as logical formulas and invoking an SMT solver over these formulas. We built a type checker for Dminor by almost directly translating the formal inference rules used to describe the bidirectional Dminor type system. In fact, we programmed so closely to the formalism that debugging an infinite loop in our implementation helped us, along with the Dminor authors, uncover a subtle typo in the formal presentation! Our Dminor type checker is 1.2K lines of Formulog. The implementation of Bierman et al. is 3.2K lines of F[#] and 400 lines of SMT-LIB; we estimate that the functionality we implemented accounts for over two thousand of these lines.⁸

The encoding of Dminor types and expressions is complex, requiring uninterpreted sorts, uninterpreted functions, universally quantified axioms, and arrays (among other features). The fact that we were able to code this relatively concisely speaks to the expressiveness of Formulog’s formula language. For example, Figure 8 shows an axiom describing the denotation of the base case of a Dminor accumulate expression, which is essentially a fold over a multiset. Here, the type `closure` is an uninterpreted sort, `enc_val` is an algebraic data type that represents an encoded Dminor value, and `accum` and `v_zero` are uninterpreted functions, where the latter represents an empty multiset.

We defined a set of mutually-recursive functions that encode expressions, environments, and types. For example, the type encoding function (fragment, Figure 9) takes a type τ and an (encoded) Dminor value v , and returns two propositions. The first is true when v has type τ . The second is a conjunction of axioms: new axioms are created to describe the denotation of the bodies of accumulate expressions as they are encountered when encoding expressions. The first case in the figure encodes the fact that any value has type `Any`. The second one says that a value has type `Bool` if it is constructed using the constructor `ev_bool`; the constructor `#is_ev_bool` is an automatically-generated constructor tester. The third case handles multiset types. It creates a fresh encoded value x , uses x to recursively create a proposition representing the encoding of the type S of items in the multiset, and then returns a proposition requiring the value to be a “good” collection (defined using the uninterpreted function `good_c`) and every item in the multiset to have type S (where `mem` is another uninterpreted function).

Although we use ML-style functions to define the logical denotation of expressions, environments, and types, we use logic programming rules to define the bidirectional type checker, which allows us to write rules that are very similar to the inference rules given in the paper. Figure 10 gives the one rule defining the subtype relation: τ is a subtype of τ_1 in environment Env if τ_1 is well formed

⁸The reference implementation is closed source; the authors have kindly provided us with line counts for each file.

```

fun encode_type(T: typ, V: enc_val smt) : bool smt * bool smt =
  match T with
  | t_any => (`true`, `true`)
  | t_bool => (`#is_ev_bool(V)`, `true`)
  | t_coll(S) =>
    let X = #{(S, V)}[enc_val] in
    let (Phi, Ax) = encode_type(S, `X`) in
    (`good_c(V) /\ forall X : mem(X, V). mem(X, V) ==> Phi`, Ax)

```

Fig. 9. This function (fragment) constructs a formula capturing the logical denotation of a Dminor type.

```

subtype(Env, T, T1) :-
  type_wf(Env, T1),
  encode_env(Env) = Phi_env,
  X = `#{(Env, T, T1)}[enc_val]`,
  encode_type(T, X) = (Phi_t, Axioms1),
  encode_type(T1, X) = (Phi_t1, Axioms2),
  Premises = [Phi_t, Phi_env, Axioms2, Axioms1, axiomatization],
  is_sat_opt(`~Phi_t1` :: Premises, z3_timeout) = some(false).

```

Fig. 10. This rule defines Dminor’s semantic subtyping relation. It uses the operator `is_sat_opt` instead of `is_valid` because its SMT queries can sometimes result in “unknown.”

and the denotation of T , given our axioms and the denotation of Env , implies the denotation of $T1$. This rule is an almost exact translation of the inference rule given in the paper.

Finally, the type checker needs to ensure that any expressions that occur in refinements are pure (i.e., terminate and are deterministic). We have written a termination checker based on the size-change principle [Lee et al. 2001]. Our implementation is another good example of the synergy between ML-style functions and Datalog rules, as we use the former to define the composition of two size-change graphs and use the latter to find the fixed point of composing size-change graphs.

We tested our type checker on six of the sample programs included in the Dminor documentation (the other three examples make use of a feature — the ability to generate an instance of a type — that we did not implement, although it should be possible to do so; to the best of our knowledge, these are the only publicly available Dminor programs). We combined these examples into a single aggregate program of ~150 LOC. The reference implementation type checked this program in 1.5 seconds using an optimization that tries syntactic subtyping before semantic subtyping; with this optimization disabled, it took 3.6 seconds.⁹ Our implementation completed in 5 seconds; it did not use this optimization (which is not detailed in the paper), but did use a newer version of Z3. Thanks to parallelization, our implementation automatically scaled to larger programs: On a synthetic program consisting of ten copies of the original aggregate program, it completed in 21.4 seconds (2.1 seconds per program copy); on a synthetic program consisting of 100 copies, it completed in 183.9 seconds (1.8 seconds per program copy). In contrast, the reference implementation did not scale: even with the syntactic-subtyping optimization enabled, it took 68 seconds on the ten-copy program and over 100 minutes on the 100-copy program.

⁹Here we used a machine with Microsoft Windows Server 2019 and the same hardware specs as our Ubuntu machine.


```

instantiate_ptsto(C, O1, Phi1, O2, widen(C, Phi_all)) :-
  instantiate_loc(C, heap(O1), heap(O2), Phi2),
  instantiate_constraint(C, Phi1, Phi3),
  Phi_all = conjoin(Phi2, Phi3).

```

Fig. 11. This rule describes how a points-to edge to object O1 labeled with constraint Phi1 is instantiated at a call site C: if at C a heap location heap(O1) can be instantiated to a heap location heap(O2) under constraint Phi2, and the original constraint on the edge Phi1 can be instantiated to a constraint Phi3, then the points-to edge to O1 labeled with Phi1 instantiates to a points-to edge to O2 labeled with widen(C, Phi_all), where Phi_all is the conjunction of Phi2 and Phi3 and widen is a function that widens constraints in mutually-recursive functions (one of the heuristics we borrowed from Scuba).

5.3 Bottom-up points-to analysis

We have implemented the bottom-up context-sensitive points-to analysis for Java proposed by Feng et al. [2015]. A points-to analysis computes a static approximation of the objects that stack variables and heap locations can point to at runtime. A bottom-up points-to analysis does this through constructing method summaries that describe the effect of a method on the heap; it is bottom-up in the sense that summaries are propagated up the call graph, from callees to callers.

In Feng et al.’s algorithm, a method summary is an abstract heap that maps abstract locations to heap objects, where an abstract location might be a stack variable, an explicitly allocated heap object, or an argument-derived heap location. Edges in the abstract heap are labeled with logical formulas that describe the conditions under which the edges hold; when a method summary is instantiated at a call site, a constraint solver can be used to filter out edges with unsatisfiable labels.

Feng et al.’s tool based on this algorithm, Scuba, is ~15K lines of Java, builds on the Chord analysis framework [Naik 2011], and uses Z3 to discharge constraints. As for many realistic static analysis tools, there is a gap between what is implemented in Scuba and the formal specification of the analysis. This is partly because Scuba is written in Java: Object-oriented programming does not naturally capture inference rules, the form of the specification. In contrast, our Formulog implementation, which is ~1.5K LOC, closely mirrors the inference rules. For example, we can directly state how a points-to edge is instantiated at a call site (Figure 11), one step of summary instantiation, a complex process defined through half a dozen mutually recursive relations that need to be computed as a fixed point. The Java code for encoding this logic is more complex and further from the formal specification. Programming close to the specification also helps check the specifications’ correctness: while implementing in Formulog one of the judgments specified by Feng et al., we discovered an inconsistency between the judgment’s definition and its type signature.

Scuba employs a range of sophisticated heuristics that are essential to making the algorithm perform in practice, as they tune precision to achieve scalability. Some go far beyond the algorithm described in the paper and are interesting in their own right. Our implementation uses some heuristics based on the ones in Scuba. The fact that we were able to implement useful heuristics — a necessity for a realistic static analysis tool — argues for the practicality of Formulog. Moreover, we were able to do so such that our code still closely reflects the core algorithm specified in the paper.

We ran both tools on the benchmarks used in the evaluation by Feng et al., which represent a selection from the pjbench suite plus the benchmark polyglot.¹⁰ These experiments include library code and use a context-sensitivity of two call sites; reflection is ignored, as are many native methods. Given an hour timeout, our implementation completed on eight of the ten benchmarks, with times ranging from five to 18 minutes (Table 1). In the median, we were 6.7× slower than Scuba. However,

¹⁰The pjbench suite is available at <https://bitbucket.org/psl-lab/pjbench/src/master/>.

Benchmark	Scuba		Formulog	
	Time	# main edges	Time	# main edges
antlr	1:11	3,313	12:16	112,415
avrora	1:05	714	7:40	127,535
hedc	0:57	867	5:04	2,962
hsqldb	0:51	780	4:53	7,039
luindex	1:40	3,395	T/O	-
polyglot	0:55	117	4:52	4,245
sunflow	3:48	7,456	T/O	-
toba-s	0:58	521	4:57	12,284
weblech	1:10	1,262	17:58	6,785
xalan	0:54	183	5:40	55,722

Table 1. In the median, our implementation of a bottom-up points-to analysis for Java was 6.7 \times slower than Scuba, the reference implementation; however, the two tools use different heuristics and thus compute very different things, as indicated by the discrepancy in the number of points-to edges computed in the summary for main (which also captures the effect on the heap of methods invoked transitively from it).

a performance comparison between the tools should be taken with a grain of salt: Since they use different heuristics, they compute very different things.

In sum, we were able to implement the algorithm in a way that is still very close to its specification and achieve decent performance on many realistic benchmarks while implementing only a small selection of heuristics. Other heuristics might have helped our version complete on the two benchmarks it timed out on. Making the algorithm practical is a significant engineering challenge: Even with its sophisticated heuristics, Scuba does not complete on all benchmarks in `pjbench`.¹¹

Moreover, our implementation could be used as a platform for exploring potential optimizations to Scuba. First, because it is automatically parallelized (with a user-chosen number of worker threads), it could be used to evaluate how well the underlying points-to algorithm parallelizes before going through all the trouble of parallelizing Scuba, which uses mutable state in a complex way. Second, thanks to the magic set transformation, we have automatically derived a goal-directed version of the analysis that computes only the summaries necessary for constructing user-requested summaries. The points-to algorithm resulting from this transformation could be used as a road map for implementing a demand-driven version of Scuba, which Feng et al. describe as future work.

5.4 Bounded symbolic evaluation

We have written a symbolic evaluator (~800 LOC) for a fragment of LLVM bitcode [Lattner and Adve 2004] corresponding to a simple imperative language with integer arrays and symbolic integers (a symbolic integer represents a set of integer values that might occur at runtime). It implements a form of bounded symbolic execution [King 1976], exploring all feasible program paths up to a given length, evaluating concretely whenever possible, and aggressively pruning infeasible paths.

Our implementation uses a different logic rule to define each of the possible cases during evaluation, and uses ML functions to manipulate and reason about complex terms representing evaluator state. For example, one rule defines when an assertion fails (Figure 12). This rule says that the path `Path` ends in a failure with evaluator state `St` if: (1) there is an `assert` instruction `Instr` with argument `X`, (2) following `Path` has led the evaluator to that instruction with state `St`, (3) `X` could have the (possibly symbolic) integer value `V` in state `St`, and (4) `V` may be zero. The function

¹¹For example, we found it timed out on `batik`, `chart`, `fop`, `lusearch`, and `pmd` (as did our implementation).

```

failed_assert(Path, St) :-
  assert_instruction(Instr, X),
  stepped(Instr, St, _, Path),
  has_value(X, St, v_int(V)),
  may_be_zero(V, St) = true.

```

Fig. 12. This rule states that the symbolic evaluator has reached a failing assertion when the argument of the assert instruction may be zero.

```

a := array of N symbolic ints;
b := symbolic int;
if (b) { sort a; }
else { sort a; assert a sorted; }

```

Fig. 13. This pseudocode sketches a C program that creates an array, branches, sorts it in each branch, but only asserts that the result is sorted in one branch.

Benchmark	# paths	KLEE	Formulog	Benchmark	# paths	KLEE	Formulog
shuffle-4	125	0:08	0:04 ($\uparrow 2.0\times$)	sort-7	22,070	25:45	3:48 ($\uparrow 6.8\times$)
shuffle-5	1,296	2:53	0:19 ($\uparrow 9.1\times$)	numbrix-sat	1	0:14	1:19 ($\downarrow 5.6\times$)
sort-6	2,718	2:25	0:18 ($\uparrow 8.1\times$)	numbrix-unsat	1	0:09	1:08 ($\downarrow 7.6\times$)

Table 2. We report absolute times for KLEE and a Formulog-based symbolic evaluation tool on six benchmark programs; for the latter, we also report speedups (\uparrow) and slowdowns (\downarrow) relative to KLEE.

`may_be_zero(V, St)` returns true if and only if V may be zero given St . We represent symbolic values as SMT formulas, so when V is symbolic, this function invokes the SMT solver.

We have evaluated our symbolic evaluator on six benchmarks based on three template programs. The first template (*shuffle- N*) non-deterministically shuffles an array of size N and asserts that the resulting array represents the same set as the input array. The second template (*sort- N* ; Figure 13) splits into two branches, sorts an array using selection sort in both branches, and asserts that the resulting array is sorted in the second branch. The third template completes a partially filled-in 4×4 grid of integers, such that there is a path from 1 to 16 where each integer follows its predecessor and only horizontal and vertical movements are used; the benchmark *numbrix-sat* runs this program on a satisfiable instance, while the benchmark *numbrix-unsat* runs it on an unsatisfiable one.

We compared our times on these benchmarks against KLEE (v2.1) [Cadaru et al. 2008], a well-known symbolic execution tool (Table 2). On programs with a single path to explore, KLEE has speedups over our version ranging from $5.6\times$ (*numbrix-sat*) to $7.6\times$ (*numbrix-unsat*). However, thanks to its automatic parallelization, Formulog scales better as the number of paths increases: For programs with more than 1000 paths, our version achieves speedups over KLEE ranging from $6.8\times$ (*sort-7*) to $9.1\times$ (*shuffle-5*). While the comparison is not totally apples-to-apples (for example, KLEE handles all of LLVM and needs to maintain a more complex symbolic state), we are encouraged by the fact that the Formulog-based version seems to scale automatically in a way that KLEE does not.

Furthermore, the Formulog-based symbolic evaluator can be run in a goal-directed mode: If we only want to check that no assertion fails, we can add the query `failed_assert(_Path, _St)`, triggering the Formulog runtime to rewrite our evaluator to explore only paths that could potentially lead to a failed assertion. On the sorting benchmarks (Figure 13), this leads to substantial speedups over KLEE — $10.4\times$ for *sort-6* and $11.8\times$ for *sort-7* — as the symbolic evaluator can ignore the first branch of the program. To further evaluate the effectiveness of this optimization, we compared our implementation against the bounded model checker CBMC (v5.11) [Clarke et al. 2004], which uses program slicing [Weiser 1984] to achieve a similar effect. On *sort-6* and *sort-7*, CBMC has speedups over our symbolic evaluator of $1.3\times$ and $1.4\times$, respectively. Thus, our automatic optimizations can put Formulog-based analyses within a reasonable factor of hand-optimized systems.

6 DESIGN EVALUATION

In this section, we evaluate the design of Formulog with respect to our case studies. We argue that Formulog is an effective and usable tool for writing SMT-based analyses.

Formulog makes it possible to write SMT-based analyses in a way that is close to their mathematical specification, leading to concise encodings. Our implementations of the Dminor type checker (Section 5.2) and the bottom-up points-to analysis (Section 5.3) directly mirror their published formal specifications; our third case study (Section 5.4), which was not based on any particular formalization, would itself be the basis of a reasonable specification of symbolic evaluation. Formulog provides language features that are a good match for the way that SMT-based analyses are specified: algebraic data types naturally encode BNF grammars (a common feature in analysis specifications); Horn clauses match judgments; ML functions fit helper functions; and the reification of formulas as terms captures the way that formulas are treated in analysis specifications.

As a corollary, analyses written in Formulog can be concise. Despite encoding quite complex logic, each of our case studies is less than 1.5K lines of code. In the case of the points-to analysis, this is 10× smaller than the reference implementation (which also uses functionality defined externally in Chord). This is partly because Scuba implements heuristics that we do not and Java is a verbose language; however, we suggest that much of the difference is because Formulog is a better fit for encoding the logic of the analysis than an imperative, object-oriented language like Java. The relative concision of Formulog matches the results reported by previous work on Datalog-based static analysis, which found that Datalog-based analyses can be orders of magnitude more concise than counterparts written in more traditional languages [Whaley et al. 2005]. The ML fragment of Formulog also helps it be concise, since ML expressions — through supporting sequenced, nested, and scoped computation — can encode logic that would be more verbose to write in Datalog.

We have shown that three diverse case studies can be naturally encoded in Formulog, suggesting that its design is a good match for a range of SMT-based analyses. However, not all analysis logic can be easily encoded in Formulog. There is currently no way to join facts, a useful operation for abstract interpretation-based analyses [Cousot and Cousot 1977]. The restriction to stratified negation is sometimes too severe: For example, one Dminor rule for the type synthesis relation *synth* is not directly expressible in Formulog, because it is defined in terms of the negation of the type well formedness relation, which is in turn defined using the relation *synth*.¹² Finally, given its lack of mutable state, Formulog is probably not a good fit for analyses that can most naturally be specified in an imperative manner, such as lazy abstraction model checking [Henzinger et al. 2002].

Because Formulog is designed to be compatible with Datalog, we can expand the type of analysis logic it supports by taking advantage of research on Datalog extensions. For instance, lattice-based recursive aggregation [Madsen et al. 2016; Szabó et al. 2018] would make it possible to join facts, and local stratification [Przymusiński 1988] would support the Dminor logic we previously cited.

Formulog provides a rich and flexible language of formulas that supports the type of logic-based reasoning found in SMT-based analyses. The formula fragment of Formulog makes it possible to use formulas the way they need to be used by static analyses. A good example of this is the decision to reify logical formulas as terms, a departure from the approach taken by constraint logic programming [Jaffar and Lassez 1987; Jaffar and Maher 1994] and constrained Horn clause (CHC) solving [Bjørner et al. 2015; Grebenschchikov et al. 2012; Gurfinkel et al. 2015; Hoder and Bjørner 2012], the two major previous paradigms for combining logic programming and constraint solving. In these systems, constraints are represented as predicates, not terms, and an inference is made if the constraints in the body of a rule are satisfiable. This approach makes sense in the

¹²To get around this, our implementation uses a less precise rule that drops the negated premise.

context of programming with *constraints*; however, it seems overly restrictive in the context of programming with *formulas*, which do not necessarily have to be used directly as constraints. For example, analyses like our Dminor type checker need to check the validity of a formula, which is the unsatisfiability of its negation. Checking validity does not easily fit in constraint-based paradigms, since constraint programming is built around satisfiability. Similarly, we might want to write an analysis that uses Craig interpolants [Craig 1957]. One could imagine extending Formulog’s SMT interface to include an operator `interpolate` that takes two formulas and returns a third (optional) formula, the interpolant; it is not clear how to do this in one of the constraint-based paradigms.

Our treatment of formula variables through the constructor $\#\{e\}[t]$ provides further evidence. This mechanism makes it easy to identify a formula variable with an object-level construct (e.g., a variable in the input program) by choosing for e the expression representing that construct. It also makes it easy to create a variable that is guaranteed to be fresh relative to a set of constructs (e.g., fresh with respect to an environment), an extremely useful operation. This is done by choosing for e a tuple of the constructs that the variable needs to be fresh with respect to. We use this trick in both the Dminor type checker and the symbolic evaluator. Crucially, this freshness mechanism is deterministic, which means that we can safely rewrite Formulog programs and parallelize them. The logic programming language Calypso [Aiken et al. 2007; Hackett 2010] provides a similar mechanism, except that it requires that all the variables in a formula are identified by terms with the same type; this severely limits its usability and is too restrictive for our case studies.

Our case studies exercise a range of the SMT-LIB standard and demonstrate the richness of our formula language. The case studies variously use algebraic data types and uninterpreted functions (the Dminor type checker and the bottom-up points-to analysis); bit vectors and arrays (the Dminor type checker and the symbolic evaluator); and integers, uninterpreted sorts, and quantifiers (the Dminor type checker). It is easy to extend Formulog with additional theories (by adding new constructors) and different types of logical reasoning (by adding new operators, like `interpolate`). As Formulog so loosely couples Datalog evaluation and constraint solving, it is easy to swap in new solver backends without major changes to the Formulog runtime; our prototype currently supports Z3 [de Moura and Bjørner 2008], CVC4 [Barrett et al. 2011], and Yices 2 [Dutertre 2014].

The design of Formulog makes it possible to advantageously apply Datalog-style optimizations to SMT-based analyses, with the result that Formulog programs can compete with analyses written in more mature languages. All of our case study implementations benefit from automatic parallelization: this scales our Dminor type checker and symbolic evaluator over the reference implementations, and helps our bottom-up points-to analysis be reasonably performant. The points-to analysis and symbolic evaluator also demonstrate the potential of the magic set transformation, as we have used it to derive demand-driven versions of these SMT-based analyses. While these types of optimizations could be added by hand to the reference implementations we compare against, the point is that the design of Formulog means that Formulog-based analyses get these optimizations for free, without the explicit effort of the analysis designer. Moreover, because of Formulog’s close affinity to Datalog, a Formulog runtime can be augmented with additional Datalog-style optimizations. For instance, a Formulog runtime could use an incremental Datalog evaluation algorithm, which efficiently evaluates Datalog programs while facts are added or retracted from EDB relations [Gupta et al. 1993; Szabó et al. 2018]. This would be helpful for using SMT-based analyses in situations where the code under analysis changes, such as in IDEs or rapidly evolving codebases.

It speaks to the design of Formulog that the high-level optimizations it enables can, in many cases, make up for the naivety of our prototype runtime. Nonetheless, we are optimistic that significantly better performance can be achieved with a sophisticated backend. As we have designed Formulog to be close to Datalog, we can take advantage of many of the optimizations that have helped Datalog

systems scale. For example, since we maintain the range restriction (which entails that every derived fact is ground), we can use concurrent data structures specialized for Datalog evaluation [Jordan et al. 2019]; since Formulog can be evaluated using standard semi-naive evaluation, we can compile Formulog programs to C++ following Soufflé’s strategy [Jordan et al. 2016].

The ML fragment is an integral part of Formulog and has a substantial impact on its usability. As discussed in Section 3.2, the first-order fragment of ML we use can be translated in a pretty straightforward way to Datalog rules, and hence can be thought of as syntactic sugar. Despite this, the ML fragment is an integral part of the Formulog programming experience. First, it improves the ergonomics of Formulog, by making it more natural to manipulate complex terms. In particular, pattern matching and let expressions provide a structured way to reflect on complex terms and sequence computation on them; this same effect is not always as easy to achieve in Datalog rules. Second, it helps Formulog achieve its design goal of allowing SMT-based analyses to be implemented in a style close to their specification, since formal specifications often involve functions in addition to inference rules. Third, it improves the performance of Formulog, as there is more overhead involved with evaluating Datalog rules than evaluating an ML expression. A substantial amount of our case study code is in the ML fragment: To give rough estimates, 65% of the Dminor type checker, 30% of the bottom-up points-to analysis, and 40% of the symbolic evaluator are in the ML fragment. Typically, the case studies use Horn clauses to define the overall structure of the analysis, and ML functions for structuring lower-level control flow, mirroring the use of judgments and helper functions in analysis specifications.

The limitation to first-order ML has several advantages. From a theoretical perspective, it means that there is an easy translation from it to Datalog rules, which allows us to give the standard Herbrand model-based semantics to Formulog programs. From a practical perspective, it ensures that we never have to unify functions, which would require higher-order unification. The specifications of our case studies did not make heavy use of higher-order functions, so they were not much missed. However, a future version of Formulog could allow a limited use of higher-order functions (for example, those programs that can be compiled to the first-order fragment).

7 RELATED WORK

Datalog-based frameworks and domain-specific languages for static analysis. A variety of static analysis frameworks have been developed based on more-or-less standard Datalog, such as bddb-ddb [Whaley et al. 2005], Chord [Naik 2011], Doop [Bravenboer and Smaragdakis 2009], QL [Av-gustinov et al. 2016], and Soufflé [Scholz et al. 2016]. Recent work has explored synthesizing Datalog-based analyses [Albarghouthi et al. 2017; Raghothaman et al. 2019]. Flix [Madsen et al. 2016] and InCA [Szabó et al. 2018] extend Datalog for analyses that operate over lattices besides the powerset lattice. InCA supports incremental evaluation, while Flix (like Formulog) includes algebraic data types and a pure functional language. Dataflow analysis is used as a case study for Datafun, a language combining Datalog and higher-order functional programming [Arntzenius and Krishnaswami 2016]. It might be possible to encode something like Formulog in Datafun; however, although it has recently been shown that Datafun can be evaluated using semi-naive evaluation [Arntzenius and Krishnaswami 2020], it is not clear to what extent other Datalog optimizations can be applied to Datafun programs. By combining Datalog with functional programming, Formulog, Flix, and Datafun are related to functional logic programming [Antoy and Hanus 2010]. The functional fragment of Formulog is less expressive than what is typically found in such languages, as Formulog functions are not first-class values and not higher-order.

Logic programming with constraints and formulas. The two dominant prior paradigms for combining logic programming and constraint solving are constraint logic programming (CLP) [Jaffar and

Lassez 1987; Jaffar and Maher 1994] and constrained Horn clause (CHC) solving [Bjørner et al. 2015; Grebenshchikov et al. 2012; Gurfinkel et al. 2015; Hoder and Bjørner 2012]. As discussed in Section 6, these systems typically encode constraints as predicates, not terms, and thus support programming with constraints as things to be satisfied, rather than programming with formulas, which can be manipulated in more interesting ways (e.g., validity checking). In the context of static analysis, these systems have been used primarily for model checking, where a model of the input system is encoded using Horn clauses [Bjørner et al. 2015; Delzanno and Podelski 1999; Flanagan 2004; Fribourg and Richardson 1996; Grebenshchikov et al. 2012]. The rules depend on the program being analyzed, and the solutions to these rules reveal properties of the model; e.g., SeaHorn [Gurfinkel et al. 2015] checks programs by solving a CHC representation of their verification conditions. This differs than the approach taken in this paper, where the rules encode an analysis independent of the input program. The Datalog mode of μZ [Hoder et al. 2011] can be thought of as a bottom-up CLP system with special support for abstract interpretation.

A few existing logic programming systems support programming with formulas (vs constraints); we would argue that none do so with the same richness and flexibility as Formulog. Codish et al. [2008] extend Prolog with an interface to a SAT solver. SICStus Prolog [Carlsson and Mildner 2012], with its CLP extensions, has been used to write model checkers [Delzanno and Podelski 1999; Fribourg and Richardson 1996; Grebenshchikov et al. 2012; Podelski and Rybalchenko 2007]; these implementations typically rely on Prolog’s non-logical features, like `assert`, making it harder to apply high-level optimizations like parallelization. Calypso [Aiken et al. 2007; Hackett 2010] is a Datalog variant that interfaces with external constraint solvers and has specialized support for bottom-up analyses. Calypso has been used with SAT and integer constraint solvers; in theory, it could be connected to an SMT solver. However, Formulog offers several advantages over Calypso for SMT-based analyses. First, Formulog’s approach to constructing formulas (via complex terms) and manipulating them (via its ML fragment) scales to the complex and heterogeneous formulas that arise in the SMT context, whereas Calypso’s approach to formulas (opaque terms, constructed via predicates) would be cumbersome in this setting. Second, Formulog’s type system supports the construction of expressive (and safe) formulas involving user-defined terms such as algebraic data types and uninterpreted functions. Third, the ML fragment of Formulog goes a long way towards making it practical for SMT-based analyses, by closing the gap between specification and implementation, and improving ergonomics and performance.

The logic programming language λ Prolog provides a natural way to represent logical formulas using a form of higher-order abstract syntax based on λ -terms and higher-order unification [Miller and Nadathur 1987; Pfenning and Elliott 1988]. Although this representation simplifies some aspects of using formulas, moving to a higher-order setting would complicate Formulog, widen the gap between Formulog and other Datalog variants, and potentially be an impediment to building a performant and scalable Formulog implementation. Answer set programming (ASP) uses specialized solvers to find a *stable model* (if it exists) of a set of Horn clauses [Brewka et al. 2011; Gelfond and Lifschitz 1988]. Common extensions support constraints on the shape of the stable model that will be found. ASP enables concise encoding of classic NP-complete constraint problems such as graph k -coloring, but it is not obviously applicable to static analysis problems.

Type system engineering. PLT Redex [Felleisen et al. 2009] and Spoofox [Kats and Visser 2010] support exploratory type system engineering. PLT Redex supports a notion of judgment modeled explicitly on inference rules. Spoofox’s type engineering framework, Statix, uses a logic programming syntax to specify type systems, with a custom solver for resolving the binding information in scope graphs simultaneously with solving typing constraints [van Antwerpen et al. 2018]. Both of these

systems use custom approaches to finding typing derivations; neither supports SMT queries, but Statix’s custom solver can resolve constraint systems that might not always terminate in Formulog.

Solver-aided languages. Scala^{Z3} [Köksal et al. 2011] supports mixed computations combining normal Scala evaluation and Z3 solving; we avoid this level of integration. Smten [Uhler and Dave 2013] is a solver-aided language that supports both concrete and symbolic evaluation; Rosette [Torlak and Bodik 2013] is a framework for creating solver-aided languages that have this property.

8 CONCLUSION

Formulog is a domain-specific language for writing SMT-based static analyses that judiciously combines Datalog, ML, and SMT solving (via an external SMT solver). As demonstrated by our case studies, it makes it possible to concisely implement a range of SMT-based analyses — refinement type checking, bottom-up points-to analysis, and symbolic evaluation — in a way close to their formal specifications, while also making it possible to automatically and advantageously apply high-level optimizations to these analyses like parallelization and goal-directed rewriting.

ACKNOWLEDGMENTS

This material is based upon work supported by the National Science Foundation under Grant No. nnnnnnn and Grant No. mmmmmmm. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- Alex Aiken, Suhabe Bugrara, Isil Dillig, Thomas Dillig, Brian Hackett, and Peter Hawkins. 2007. An Overview of the Saturn Project. In *Proceedings of the 7th ACM SIGPLAN-SIGSOFT Workshop on Program Analysis for Software Tools and Engineering*. 43–48.
- Aws Albarghouthi, Paraschos Koutris, Mayur Naik, and Calvin Smith. 2017. Constraint-Based Synthesis of Datalog Programs. In *Proceedings of the 23rd International Conference on Principles and Practice of Constraint Programming*. 689–706.
- Sergio Antoy and Michael Hanus. 2010. Functional Logic Programming. *Commun. ACM* 53, 4 (2010), 74–85.
- Krzysztof R Apt, Howard A Blair, and Adrian Walker. 1988. Towards a Theory of Declarative Knowledge. In *Foundations of Deductive Databases and Logic Programming*. Elsevier, 89–148.
- Molham Aref, Balder ten Cate, Todd J Green, Benny Kimelfeld, Dan Olteanu, Emir Pasalic, Todd L Veldhuizen, and Geoffrey Washburn. 2015. Design and Implementation of the LogicBlox System. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data*. 1371–1382.
- Michael Arntzenius and Neel Krishnaswami. 2020. Seminäive Evaluation for a Higher-Order Functional Language. *Proceedings of the ACM on Programming Languages* 4, POPL (2020), 22:1–22:28.
- Michael Arntzenius and Neelakantan R. Krishnaswami. 2016. Datafun: A Functional Datalog. In *Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming*. 214–227.
- Pavel Avgustinov, Oege De Moor, Michael Peyton Jones, and Max Schäfer. 2016. QL: Object-Oriented Queries on Relational Data. In *Proceedings of the 30th European Conference on Object-Oriented Programming*. 2:1–2:25.
- Francois Bancilhon. 1986. Naive Evaluation of Recursively Defined Relations. In *On Knowledge Base Management Systems*. Springer, 165–178.
- Francois Bancilhon, David Maier, Yehoshua Sagiv, and Jeffrey D Ullman. 1985. Magic Sets and Other Strange Ways to Implement Logic Programs. In *Proceedings of the Fifth ACM SIGACT-SIGMOD Symposium on Principles of Database Systems*. 1–15.
- Clark Barrett, Christopher L. Conway, Morgan Deters, Liana Hadarean, Dejan Jovanović, Tim King, Andrew Reynolds, and Cesare Tinelli. 2011. CVC4. In *Proceedings of the 23rd International Conference on Computer Aided Verification*. 171–177.
- Clark Barrett, Pascal Fontaine, and Cesare Tinelli. 2017. *The SMT-LIB Standard: Version 2.6*. Technical Report. Department of Computer Science, The University of Iowa.
- Catriel Beeri and Raghu Ramakrishnan. 1991. On the Power of Magic. *The Journal of Logic Programming* 10, 3-4 (1991), 255–299.
- Gavin M. Bierman, Andrew D. Gordon, Cătălin Hrițcu, and David Langworthy. 2012. Semantic Subtyping with an SMT Solver. *Journal of Functional Programming* 22, 1 (2012), 31–105.

- Nikolaj Bjørner, Arie Gurfinkel, Ken McMillan, and Andrey Rybalchenko. 2015. Horn Clause Solvers for Program Verification. In *Fields of Logic and Computation II*. Springer, 24–51.
- Martin Bravenboer and Yannis Smaragdakis. 2009. Strictly Declarative Specification of Sophisticated Points-to Analyses. In *Proceedings of the 24th ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*. 243–262.
- Gerhard Brewka, Thomas Eiter, and Mirosław Truszczyński. 2011. Answer Set Programming at a Glance. *Commun. ACM* 54, 12 (2011), 92–103.
- Cristian Cadar, Daniel Dunbar, and Dawson Engler. 2008. KLEE: Unassisted and Automatic Generation of High-Coverage Tests for Complex Systems Programs. In *Proceedings of the 8th USENIX Conference on Operating Systems Design and Implementation*. 209–224.
- Cristian Cadar and Koushik Sen. 2013. Symbolic Execution for Software Testing: Three Decades Later. *Commun. ACM* 56, 2 (Feb. 2013), 82–90.
- Mats Carlsson and Per Mildner. 2012. SICStus Prolog—The First 25 years. *Theory and Practice of Logic Programming* 12, 1-2 (2012), 35–66.
- Alessandro Cimatti and Alberto Griggio. 2012. Software Model Checking via IC3. In *Proceedings of the 24th International Conference on Computer Aided Verification*. 277–293.
- Edmund Clarke, Daniel Kroening, and Flavio Lerda. 2004. A Tool for Checking ANSI-C Programs. In *Proceedings of the 10th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. 168–176.
- Michael Codish, Vitaly Lagoon, and Peter J Stuckey. 2008. Logic Programming with Satisfiability. *Theory and Practice of Logic Programming* 8, 1 (2008), 121–128.
- Patrick Cousot and Radhia Cousot. 1977. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Proceedings of the 4th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. 238–252.
- William Craig. 1957. Three Uses of the Herbrand-Gentzen Theorem in Relating Model Theory and Proof Theory. *The Journal of Symbolic Logic* 22, 3 (1957), 269–285.
- Luis Damas and Robin Milner. 1982. Principal Type-Schemes for Functional Programs. In *Proceedings of the 9th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 207–212.
- Leonardo de Moura and Nikolaj Bjørner. 2008. Z3: An Efficient SMT Solver. In *Proceedings of the 14th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. 337–340.
- Giorgio Delzanno and Andreas Podelski. 1999. Model Checking in CLP. In *Proceedings of the 5th International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. 223–239.
- David Detlefs, Greg Nelson, and James B. Saxe. 2005. Simplify: A Theorem Prover for Program Checking. *J. ACM* 52, 3 (2005), 365–473.
- Bruno Dutertre. 2014. Yices 2.2. In *Proceedings of the 26th International Conference on Computer Aided Verification*. 737–744.
- Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. 2009. *Semantics Engineering with PLT Redex* (1st ed.). The MIT Press.
- Yu Feng, Xinyu Wang, Isil Dillig, and Thomas Dillig. 2015. Bottom-up Context-Sensitive Pointer Analysis for Java. In *Proceedings of the 13th Asian Symposium on Programming Languages and Systems*. 465–484.
- Cormac Flanagan. 2004. Automatic Software Model Checking via Constraint Logic. *Science of Computer Programming* 50, 1-3 (2004), 253–270.
- Antonio Flores-Montoya and Eric Schulte. 2019. Datalog Disassembly. arXiv:arXiv:1906.03969
- Laurent Fribourg and Julian Richardson. 1996. Symbolic Verification with Gap-Order Constraints. In *Proceedings of the 6th International Workshop on Logic Programming Synthesis and Transformation*. 20–37.
- Hervé Gallaire and Jack Minker (Eds.). 1978. *Logic and Data Bases*. Plenum Press.
- Michael Gelfond and Vladimir Lifschitz. 1988. The Stable Model Semantics for Logic Programming. In *Proceedings of the 5th International Conference and Symposium on Logic Programming*. 1070–1080.
- Sergey Grebenshchikov, Nuno Lopes, Corneliu Popeea, and Andrey Rybalchenko. 2012. Synthesizing Software Verifiers from Proof Rules. In *Proceedings of the 33rd ACM SIGPLAN Conference on Programming Language Design and Implementation*. 405–416.
- Neville Grech, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. 2019. Gigahorse: Thorough, Declarative Decompilation of Smart Contracts. In *Proceedings of the 41st International Conference on Software Engineering*. 1176–1186.
- Neville Grech, Michael Kong, Anton Jurisevic, Lexi Brent, Bernhard Scholz, and Yannis Smaragdakis. 2018. Madmax: Surviving Out-of-Gas Conditions in Ethereum Smart Contracts. *Proceedings of the ACM on Programming Languages* 2, OOPSLA (2018), 116:1–116:27.
- Todd J. Green, Shan Shan Huang, Boon Thau Loo, and Wenchao Zhou. 2013. Datalog and Recursive Query Processing. *Foundations and Trends in Databases* 5, 2 (2013), 105–195.

- Salvatore Guarnieri and V Benjamin Livshits. 2009. GATEKEEPER: Mostly Static Enforcement of Security and Reliability Policies for JavaScript Code. In *Proceedings of the 18th USENIX Security Symposium*. 78–85.
- Ashish Gupta, Inderpal Singh Mumick, and Venkatramanan Siva Subrahmanian. 1993. Maintaining Views Incrementally. *ACM SIGMOD Record* 22, 2 (1993), 157–166.
- Arie Gurfinkel, Temesghen Kahsay, Anvesh Komuravelli, and Jorge A Navas. 2015. The SeaHorn Verification Framework. In *Proceedings of the 27th International Conference on Computer Aided Verification*. 343–361.
- Brian Hackett. 2010. *Type Safety in the Linux Kernel*. Ph.D. Dissertation. Stanford University.
- Thomas A. Henzinger, Ranjit Jhala, Rupak Majumdar, and Grégoire Sutre. 2002. Lazy Abstraction. In *Proceedings of the 29th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 58–70.
- Roger Hindley. 1969. The Principal Type-Scheme of an Object in Combinatory Logic. *Trans. Amer. Math. Soc.* 146 (1969), 29–60.
- Kryštof Hoder and Nikolaj Bjørner. 2012. Generalized Property Directed Reachability. In *Proceedings of the 15th International Conference on Theory and Applications of Satisfiability Testing*. Springer, 157–171.
- Kryštof Hoder, Nikolaj Bjørner, and Leonardo De Moura. 2011. μZ —An Efficient Engine for Fixed Points with Constraints. In *Proceedings of the 23rd International Conference on Computer Aided Verification*. 457–462.
- Joxan Jaffar and Jean-Louis Lassez. 1987. Constraint Logic Programming. In *Proceedings of the 14th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*. 111–119.
- Joxan Jaffar and Michael J. Maher. 1994. Constraint Logic Programming: A Survey. *The Journal of Logic Programming* 19 (1994), 503–581.
- Herbert Jordan, Bernhard Scholz, and Pavle Subotić. 2016. Soufflé: On Synthesis of Program Analyzers. In *Proceedings of the 28th International Conference on Computer Aided Verification*. 422–430.
- Herbert Jordan, Pavle Subotic, David Zhao, and Bernhard Scholz. 2019. A Specialized B-tree for Concurrent Datalog Evaluation. In *Proceedings of the 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. 327–339.
- Lennart C.L. Kats and Eelco Visser. 2010. The Spoox Language Workbench: Rules for Declarative Specification of Languages and IDEs. In *Proceedings of the 25th ACM International Conference on Object-Oriented Programming, Systems, Languages, and Applications*. 444–463.
- James C. King. 1976. Symbolic Execution and Program Testing. *Commun. ACM* 19, 7 (1976), 385–394.
- Ali Sinan Köksal, Viktor Kuncak, and Philippe Suter. 2011. Scala to the Power of Z3: Integrating SMT and Programming. In *Proceedings of the 23rd International Conference on Automated Deduction*. 400–406.
- Robert Kowalski. 1979. Algorithm = Logic + Control. *Commun. ACM* 22, 7 (1979), 424–436.
- Chris Lattner and Vikram Adve. 2004. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the 2nd IEEE/ACM International Symposium on Code Generation and Optimization*. 75–88.
- Chin Soon Lee, Neil D. Jones, and Amir M. Ben-Amram. 2001. The Size-Change Principle for Program Termination. In *Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. 81–92.
- V. Benjamin Livshits and Monica S. Lam. 2005. Finding Security Vulnerabilities in Java Applications with Static Analysis. In *Proceedings of the 14th USENIX Security Symposium*. 271–286.
- Magnus Madsen, Ming-Ho Yee, and Ondřej Lhoták. 2016. From Datalog to Flix: a Declarative Language for Fixed Points on Lattices. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 194–208.
- Kenneth L McMillan. 2006. Lazy Abstraction with Interpolants. In *Proceedings of the 18th International Conference on Computer Aided Verification*. Springer, 123–136.
- Michael Meskes and Jörg Noack. 1993. The Generalized Supplementary Magic-Sets Transformation for Stratified Datalog. *Inform. Process. Lett.* 47, 1 (1993), 31–41.
- Dale Miller and Gopalan Nadathur. 1987. A Logic Programming Approach to Manipulating Formulas and Programs. In *Proceedings of the 1987 Symposium on Logic Programming*. 379–388.
- Inderpal Singh Mumick, Hamid Pirahesh, and Raghu Ramakrishnan. 1990. The Magic of Duplicates and Aggregates. In *Proceedings of the 16th International Conference on Very Large Data Bases*. 264–277.
- Mayur Naik. 2011. Chord: A Program Analysis Platform for Java. https://www.seas.upenn.edu/~mhnaik/chord/user_guide/index.html. Accessed: 2020-04-01.
- Frank Pfenning and Conal Elliott. 1988. Higher-Order Abstract Syntax. In *Proceedings of the ACM SIGPLAN 1988 Conference on Programming Language Design and Implementation*. 199–208.
- Andreas Podelski and Andrey Rybalchenko. 2007. ARMC: The Logical Choice for Software Model Checking with Abstraction Refinement. In *Proceedings of the 9th International Symposium on Practical Aspects of Declarative Languages*. 245–259.
- Teodor C Przymusiński. 1988. On the Declarative Semantics of Deductive Databases and Logic Programs. In *Foundations of Deductive Databases and Logic Programming*. Elsevier, 193–216.

- Mukund Raghothaman, Jonathan Mendelson, David Zhao, Mayur Naik, and Bernhard Scholz. 2019. Provenance-Guided Synthesis of Datalog Programs. *Proceedings of the ACM on Programming Languages* 4, POPL (2019), 1–27.
- Thomas W. Reps. 1995. Demand Interprocedural Program Analysis Using Logic Databases. In *Proceedings of the 3rd ACM SIGSOFT Symposium on Foundations of Software Engineering*. 163–196.
- Patrick M. Rondon, Ming Kawaguci, and Ranjit Jhala. 2008. Liquid Types. In *Proceedings of the 29th ACM SIGPLAN Conference on Programming Language Design and Implementation*. 159–169.
- Bernhard Scholz, Herbert Jordan, Pavle Subotić, and Till Westmann. 2016. On Fast Large-Scale Program Analysis in Datalog. In *Proceedings of the 25th International Conference on Compiler Construction*. 196–206.
- Yannis Smaragdakis and Martin Bravenboer. 2011. Using Datalog for Fast and Easy Program Analysis. In *Datalog Reloaded*. Springer, 245–251.
- Tamás Szabó, Gábor Bergmann, Sebastian Erdweg, and Markus Voelter. 2018. Incrementalizing Lattice-Based Program Analyses in Datalog. *Proceedings of the ACM on Programming Languages* 2, OOPSLA (2018), 139:1–139:29.
- Emina Torlak and Rastislav Bodik. 2013. Growing Solver-Aided Languages with Rosette. In *Proceedings of the 2013 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming & Software*. 135–152.
- Petar Tsankov, Andrei Dan, Dana Drachler-Cohen, Arthur Gervais, Florian Buenzli, and Martin Vechev. 2018. Securify: Practical Security Analysis of Smart Contracts. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*. 67–82.
- Richard Uhler and Nirav Dave. 2013. Smten: Automatic Translation of High-Level Symbolic Computations into SMT Queries. In *Proceedings of the 25th International Conference on Computer Aided Verification*. 678–683.
- Hendrik van Antwerpen, Casper Bach Poulsen, Arjen Rouvoet, and Eelco Visser. 2018. Scopes as Types. *Proceedings of the ACM on Programming Languages* 2, OOPSLA (2018), 114:1–114:30.
- Allen Van Gelder. 1989. Negation as Failure Using Tight Derivations for General Logic Programs. *The Journal of Logic Programming* 6, 1-2 (1989), 109–133.
- Mark Weiser. 1984. Program Slicing. *IEEE Transactions on Software Engineering* 4 (1984), 352–357.
- John Whaley, Dzintars Avots, Michael Carbin, and Monica S. Lam. 2005. Using Datalog with Binary Decision Diagrams for Program Analysis. In *Proceedings of the Third Asian Symposium on Programming Languages and Systems*. 97–118.
- John Whaley and Monica S. Lam. 2004. Cloning-Based Context-Sensitive Pointer Alias Analysis Using Binary Decision Diagrams. In *Proceedings of the ACM SIGPLAN 2004 Conference on Programming Language Design and Implementation*. 131–144.

Types	
Types	$\tau ::= t \mid t \text{ smt} \mid t \text{ sym} \mid \text{model}$
Pre-types	$t ::= B \mid D \vec{\tau} \mid \alpha$
Base types	$B ::= \text{bool} \mid \text{bv}[k]_{k \in \mathbb{N}^+} \mid \dots$
Contexts	
Data type declarations	$\Delta ::= \cdot \mid \Delta, D : \forall \vec{\alpha}_i. \overrightarrow{\{c_j : \vec{\tau}_k\}}$
Program declarations	$\Phi ::= \cdot \mid \Phi, f : \forall \vec{\alpha}, \vec{\tau} \rightarrow \tau \mid \Phi, uf : \vec{\tau} \rightarrow t \mid \Phi, p \subseteq \vec{\tau}$
Variable contexts	$\Gamma ::= \cdot \mid \Gamma, x : \tau \mid \Gamma, \alpha$
Terms	
Programs	$\text{prog} ::= \vec{F}_i \vec{H}_j$
Functions	$F ::= \text{fun } f(\vec{X}_i : \vec{\tau}_i) : \tau = e$
Horn clauses	$H ::= p(\vec{X}_i) :- \vec{P}_j$
Premises	$P ::= A \mid !A$
Atoms	$A ::= p(\vec{e}_i) \mid X = e$
Expressions	$e ::= k \mid X \mid c(\vec{e}_i) \mid f(\vec{e}_i) \mid p(\vec{e}_i) \mid \otimes(\vec{e}_i) \mid \backslash \phi$ $\text{let } X = e_1 \text{ in } e_2 \mid \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \mid$ $\text{match } e \text{ with } \overrightarrow{c_i(\vec{X}_j) \rightarrow e_i}$
SMT formulas	$\phi ::= c_{\text{var}}^{\text{SMT}}[x, t]() \mid c_{\text{const}}^{\text{SMT}}[k]() \mid c_{\text{let}}^{\text{SMT}}(\phi_1, \phi_2, \phi_3) \mid$ $c_{\text{ctor}}^{\text{SMT}}[c](\vec{\phi}_i) \mid c_{\text{forall}}^{\text{SMT}}(\phi_1, \phi_2) \mid c_{\text{uf}}^{\text{SMT}}[uf](\vec{\phi}_i) \mid ,e$
Constants	$k ::= \text{true} \mid \text{false} \mid 0 \mid 1 \mid \dots$
Namespaces	
Type modes	$m ::= \text{exp} \mid \text{smt}$
Data type names	$D \in \text{ADTVar}$
Type variables	$\alpha \in \text{TVar}$
Constructors	$c \in \text{CtorVar}$
Formulog variables	$X \in \text{Var}$
SMT variables	$x \in \text{SMTVar}$
Predicates	$p \in \text{PredVar}$
Functions	$f \in \text{FunVar}$
Uninterpreted functions	$uf \in \text{UninterpFunVar}$

Fig. 14. Syntax of Formulog’s formal model

A FORMULOG’S FORMAL MODEL

We define a ‘middleweight’ formal model of Formulog, designing a type system (Section B) and an operational semantics (Section C), relating the two in a proof of type safety (Section D).

Our model characterizes Formulog as a two-level system (Figure 14), comprising Datalog-esque Horn clauses H and first-order functions F ; Horn clause “rules” are made up of premises P , where each premise is a series of (possibly negated) atoms A . Each atom A either references a Datalog predicate or binds a variable to an expression e . Expressions themselves have two mutually recursive

modes: ordinary functional computation e and quoted SMT terms ϕ , which can include unquoted expressions e .

The Datalog fragment of Formulog is fairly standard syntactically, up to the addition of the atomic form $X = e$. We constrain premises to a sort of administrative normal form: predicate references apply only to variables, written $p(\vec{X}_i)$, and expression constraints bind variables, as in $Y = e$. Our implementation can handle compound premises like $p(e_1, e_2)$; our formal model would require rewriting such a premise to three premises: $p(X, Y)$, $X = e_1$, and $Y = e_2$ (for some fresh X and Y).

The functional programming fragment fully annotates the types on its functions F ; variable names in both fragments are written in capital letters. (Our implementation merely demands that the first letter be capitalized.) SMT variables are written using lowercase letters and annotated with their type, as in $c_{\text{var}}^{\text{SMT}}[x, t]()$. As described in Section 3.3, our implementation allows any value to be used as the name of an SMT variable; here, without loss of generality, we treat SMT variables as being drawn from a distinct universe. Code in the functional fragment can treat Datalog relations as predicates, i.e., $p(\vec{v}_i)$ returns true when $\vec{v}_i \in p$. In our implementation, some elements of \vec{v}_i can be the wildcard $??$, turning a Datalog predicate into a list. For example, if $p \subseteq \text{bool} \times \text{bv}[32]$, then: $p(\text{true}, 42)$ yields a bool; $p(??, 42)$ returns a list of bools b such that $p(b, 42)$ holds; $p(\text{true}, ??)$ returns a list of $\text{bv}[32]$ s n such that $p(\text{true}, n)$ holds; $p(??, ??)$ returns a list of $\text{bool} \times \text{bv}[32]$, i.e., the relation p . We don't include this behavior in our formal model.

The set of available base types B must include bool at a minimum; any other SMT-embeddable base types are acceptable, e.g., k -width bit vectors for a statically known k .

As a matter of notation, we write \vec{e}_i for a metavariable e to mean a possibly empty sequence of e s, indexed by i . When more than one variable shares the same index, we mean that those sequences must be of the same length (e.g., in the match syntax, each branch of a match is a triple of a constructor c , a vector of variable names for c 's arguments, and a corresponding single expression).

B FORMULOG'S TYPE SYSTEM

We begin by presenting type checking rules for Formulog (Figures 15, 16, 17, and 18). Our *implementation* of Formulog not only performs type checking, but can also perform type inference, e.g., automatically finding type variable substitutions.

Our types are broken into two levels: types τ and pre-types t . Every pre-type t can be directly considered as a type, but there are two additional types: t smt, the type of SMT formulas yielding t , and t sym, the type of SMT variables of type t . We factor the syntax in this way to prevent anomalies like bool smt smt , which would mean SMT formulas that yield SMT formulas that yield booleans. It is *not* the case, however, that every pre-type t is necessarily representable as an SMT type, because data types may contain SMT formulas as arguments; we discuss how we categorize SMT-representable types shortly.

Before we begin, some further notational clarification. Rules are named by their primary subjects followed by a hyphen and a descriptive name. Whenever we use indices in rules, we will always map (stating a single premise in terms of the index, e.g., prog-WF) or fold (stating first, indexed, and last, e.g., $\vec{X}\vec{\tau}$ -ALL) over the sequence. We omit the indices when selecting an element of a sequence or set (as in, e.g., e -MATCH in Figure 17).

All of our typing rules are in terms of a fixed set of data type declarations Δ and program declarations Φ (Figure 14). Data type declarations Δ map data type names D to some number of type arguments \vec{a}_i and a set of constructors \vec{c}_j , each of which takes some number of arguments of type $\vec{\tau}_k$; each c_j can have a different number of arguments. Program declarations Φ collect the

Type and typing context well formedness

$$\begin{array}{c}
\frac{}{\vdash \cdot} \quad \Gamma\text{-EMPTY} \qquad \frac{\vdash \Gamma \quad \Gamma \vdash_{\text{exp}} \tau}{\vdash \Gamma, x : \tau} \quad \Gamma\text{-VARIABLE} \qquad \frac{\Delta \vdash \Gamma \quad \Delta; \Gamma \vdash_m \tau}{\vdash \Gamma, \alpha} \quad \Gamma\text{-TVARIABLE} \\
\\
\frac{}{\Gamma \vdash_m B} \quad t\text{-BASE} \qquad \frac{\alpha \in \Gamma}{\Gamma \vdash_{\text{exp}} \alpha} \quad t\text{-TVARIABLE} \qquad \frac{\Delta(D) = \forall \vec{\alpha}_i, \{ \dots \} \quad \Gamma \vdash_m \tau_i}{\Gamma \vdash_m D \vec{\tau}_i} \quad t\text{-ADT} \\
\\
\frac{\Gamma \vdash_{\text{smt}} t}{\Gamma \vdash_m t \text{ smt}} \quad \tau\text{-SMT} \qquad \frac{\Gamma \vdash_{\text{smt}} t}{\Gamma \vdash_m t \text{ sym}} \quad \tau\text{-SYM} \qquad \frac{}{\Gamma \vdash_{\text{exp}} \text{model}} \quad \tau\text{-MODEL}
\end{array}$$

Data type and program signature well formedness

$$\begin{array}{c}
\vdash \Delta \Leftrightarrow \forall D : \forall \vec{\alpha}. \{c_1 : \vec{\tau}_1, \dots, c_n : \vec{\tau}_n\} \in \Delta \forall i, \\
(1) \forall D' \in \text{dom}(\Delta), c_i \in \Delta(D') \Rightarrow D = D' \\
(2) \vec{\alpha} \vdash_{\text{exp}} \tau_i \\
(3) \forall \beta \in \vec{\alpha}, \beta \in \vec{\tau}_i \\
\\
\frac{}{\vdash \cdot} \quad \Phi\text{-EMPTY} \qquad \frac{\vdash \Phi \quad \forall \beta \in \vec{\alpha}_i, \beta \in \vec{\tau}_j, \tau \quad \vec{\alpha}_i \vdash_{\text{exp}} \tau_j \quad \vec{\alpha}_i \vdash_{\text{exp}} \tau}{\vdash \Phi, f : \forall \vec{\alpha}_i, \vec{\tau}_j \rightarrow \tau} \quad \Phi\text{-FUNCTION} \\
\\
\frac{\vdash \Phi \quad \cdot \vdash_{\text{exp}} \tau_i}{\vdash \Phi, p \subseteq \vec{\tau}_i} \quad \Phi\text{-REL} \qquad \frac{\vdash \Phi \quad \cdot \vdash_{\text{smt}} t_i \quad \cdot \vdash_{\text{smt}} t}{\vdash \Phi, uf : \vec{t}_i \rightarrow t} \quad \Phi\text{-UFUNCTION}
\end{array}$$

Program and function typing

$$\frac{\vdash \Delta \quad \vdash \Phi \quad \Delta; \Phi \vdash F_i \quad \Delta; \Phi \vdash H_j}{\Delta; \Phi \vdash \vec{F}_i \vec{H}_j} \quad \text{prog-WF}$$

Fig. 15. Type, context, and definition well formedness; top-level program typing

signatures of first-order polymorphic functions $f : \forall \vec{\alpha}, \vec{\tau} \rightarrow \tau$, uninterpreted functions for use in the SMT solver $uf : \vec{t} \rightarrow t$, and relations $p \subseteq \vec{\tau}_i$.

The highest level typing rule is prog-WF (Figure 15), which ensures that the declarations are well formed and each part of the program is well formed.

The context and type well formedness rules (Figure 15) are mostly straightforward, type well formedness being the most interesting. Each type can be found to be well formed in either SMT mode *smt*—i.e., it can be exported to the SMT solver—or in expression mode *exp*, meaning it cannot be. There is a sub-moding relationship: well formed types at *smt* are also well formed at *exp*, but not necessarily vice-versa: for example, there is no way to export an SMT formula or variable as the *object* of another SMT formula, only as a constituent. We assume that all Formulog constants are SMT representable, i.e., $\cdot \vdash_{\text{smt}} \text{typeof}(k)$ for all constants k . Data type declarations are polymorphic, but we disallow phantom type variables. Data types can freely mutually recurse. Uninterpreted functions must be in terms of pre-types, and those pre-types must be closed and SMT representable (*smt*); functions and relations can use any well formed types (*exp*). Functions can be polymorphic but we disallow phantom type variables; relations have monomorphic types. Disallowing phantom types in constructors and functions and keeping relations monomorphic ensure that these forms are “reverse determinate”, i.e., the types of their arguments uniquely determine their types.

Variable binding and typing

$$\boxed{\Gamma \vdash x, \tau \triangleright \Gamma}$$

$$\boxed{\Gamma \vdash \vec{x}, \vec{\tau} \triangleright \Gamma}$$

$$\frac{X \notin \text{dom}(\Gamma)}{\Gamma \vdash X, \tau \triangleright \Gamma, X:\tau} \quad X\tau\text{-BIND}$$

$$\frac{\Gamma(X) = \tau}{\Gamma \vdash X, \tau \triangleright \Gamma} \quad X\tau\text{-CHECK}$$

$$\frac{\Gamma \vdash X_0, \tau_0 \triangleright \Gamma_1 \quad \dots \quad \Gamma_i \vdash X_i, \tau_i \triangleright \Gamma_{i+1} \quad \dots \quad \Gamma_n \vdash X_n, \tau_n \triangleright \Gamma'}{\Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'} \quad \vec{X}\vec{\tau}\text{-ALL}$$

Premise typing

$$\boxed{\Delta; \Phi; \Gamma \vdash P \triangleright \Gamma}$$

$$\frac{p \subseteq \vec{\tau}_i \in \Phi \quad \Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'}{\Gamma \vdash p(\vec{X}_i) \triangleright \Gamma'} \quad P\text{-PosATOM}$$

$$\frac{p \subseteq \vec{\tau}_i \in \Phi \quad \Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma}{\Gamma \vdash !p(\vec{X}_i) \triangleright \Gamma} \quad P\text{-NEGATOM}$$

$$\frac{\vec{X}_i \not\subseteq \Gamma \quad \Delta(D) = \forall \vec{\alpha}_j, \{\dots, c : \vec{\tau}_i, \dots\} \quad \Gamma \vdash Y, D \vec{\tau}'_j \triangleright \Gamma \quad \Gamma \vdash \vec{X}_i, \vec{\tau}_i [\vec{\tau}'_j / \vec{\alpha}_j] \triangleright \Gamma'}{\Gamma \vdash Y = c(\vec{X}_i) \triangleright \Gamma'} \quad P\text{-EQCTOR-BF}$$

$$\frac{\vec{X}_i \not\subseteq \Gamma \quad \Gamma \vdash c_c^{\text{SMT}} : \vec{\tau}_i \rightarrow \tau \quad \Gamma \vdash Y, \tau \triangleright \Gamma \quad \Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'}{\Gamma \vdash Y = `c_c^{\text{SMT}}(\vec{X}_i)` \triangleright \Gamma'} \quad P\text{-EQSMT-BF}$$

$$\frac{\Gamma \vdash e : \tau \quad \Gamma \vdash Y, \tau \triangleright \Gamma'}{\Gamma \vdash Y = e \triangleright \Gamma'} \quad P\text{-EQ-FB}$$

$$\frac{\Gamma \vdash e : \tau \quad \Gamma \vdash Y, \tau \triangleright \Gamma}{\Gamma \vdash !(Y = e) \triangleright \Gamma} \quad P\text{-NEGEQ}$$

Clause typing

$$\boxed{\Delta; \Phi \vdash H}$$

$$\frac{\cdot \vdash P_0 \triangleright \Gamma_1 \quad \dots \quad \Gamma_j \vdash P_j \triangleright \Gamma_{j+1} \quad \dots \quad \Gamma_n \vdash P_n \triangleright \Gamma' \quad p \subseteq \vec{\tau}_i \in \Phi \quad \Gamma' \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'}{\vdash p(\vec{X}_i) :- \vec{P}_j} \quad H\text{-CLAUSE}$$

Fig. 16. Typing rules: Horn clauses (rules)

Since the declaration environments Δ and Φ are statically determined for an entire program, we typically leave them *implicit*. Implicit parameters are in gray in the boxed rule schemata in the figures. In proofs we will treat these parameters explicitly, but we conserve space by stating the rules without threading implicit parameters through. For example, the data type declarations Δ are necessary to ensure that *t*-ADT only allows us to name data types that have actually been defined. Rather than threading Δ through every rule for context and type well formedness, we write Δ in the rule schemata.

The type checking of the Datalog fragment of Formulog (Figure 16) must encode two Datalog invariants in addition to conventional typing constraints: the range restriction, i.e., every variable in the head of a rule appears somewhere in a premise; and appropriate binding, i.e., it is possible to interpret a Horn clause in such a way that all of the variables will be bound at the end. Our formal rules ensure that the program has correct binding structure for a left-to-right evaluation of each

Horn clause. An implementation could determine whether or not an ordering would work and could reorder programs into an appropriate order automatically. Our formal model does not enforce that the dependencies between relations are appropriately stratified, though doing so would be easy: the relation-and-function call graph should not have any “negative” edge in a cycle, where a negative edge is created whenever there is a negated predicate in a rule body or a predicate is invoked as a function.

Concretely, *H-CLAUSE* ensures that (a) a left-to-right binding order produces some appropriate final context Γ' (via the premise typing judgment), (b) the range restriction is satisfied, ($\vec{X}_i \subseteq \Gamma'$) by making sure that (c) every variable is well typed and bound ($\Gamma' \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'$ —having the same Γ' means no new bindings were introduced when checking the head variables).

Premise typing $\Gamma \vdash P \triangleright \Gamma$ and variable binding and typing $\Gamma \vdash x, \tau \triangleright \Gamma$ work together to generate appropriate types for each premise. Positive references to relations are well formed in binding Γ' according to *P-PosATOM* when (a) the use is well typed ($p \subseteq \vec{\tau}_i \in \Phi$) and (b) the variables used in the premise yield the binding Γ' . Negative references to relations $!p(\vec{X}_i)$ additionally require that all of the X_i be already bound, i.e., the resulting Γ is the same as the starting one. We split expression equality constraints $Y = e$ into three main cases:

- (1) Y is bound and e is a constructor $c(\vec{X}_i)$ where all of \vec{X}_i are unbound (*P-EQCTOR-BF*).
- (2) Y is bound and e is a quoted SMT constructor $\text{`}c_c^{\text{SMT}}(\vec{X}_i)\text{'}$ where all of \vec{X}_i are unbound (*P-EQSMT-BF*).
- (3) Y is possibly unbound and e has no unbound variables (*P-EQ-FB*).

It is critical that we avoid the case where both Y and some of the X_i are unbound, in which case we would need to perform true unification (or even higher-order unification, depending on our treatment of functional programs). In the case where Y is bound and the expression has no unbound variables, only the *P-EQ-FB* case could apply. There is a fourth, irrelevant case: *P-NEGEQ*. No binding can possibly occur there, so the constraint is simply checked by running e and making sure it isn't equal to Y .

The binding rules come in three forms: *X τ -BIND* for adding a new binding, *X τ -CHECK* for ensuring that an already bound variable is matched at appropriate type, and a vectorized form *X $\vec{\tau}$ -ALL* for folding over a sequence of such bindings. Note that the resulting bindings are the same, i.e., $\Gamma \vdash X, \tau \triangleright \Gamma$, if and only if $X \in \text{dom}(\Gamma)$; the same holds for vectors of variables and types, as well (Lemma E.5).

We split the rules for expressions e and formulas ϕ in two parts (Figures 17 and 18, respectively). Expression typing is conventional for functional languages. We adopt a declarative style for type substitutions (*e-CTOR*, *e-FUN*, *e-OP*, *e-MATCH*). Our actual implementation uses Hindley–Damas–Milner type inference [Damas and Milner 1982; Hindley 1969] to find the correct types to use. As to Formulog-specific features, we ensure type well formedness is in exp-mode; the $\text{`}\phi\text{'}$ expression switches from expression mode to formula mode. Relations $p \subseteq \vec{\tau}_i \in \Phi$ are treated as if they are functions of type $\vec{\tau}_i \rightarrow \text{bool}$. Since Datalog predicates can occur in functional terms, we must use the control-flow graph of the program when analyzing for stratification. Consider the following program:

```
fun f(X : bv[32]) : bv[32] = if p(X) then ... else X
p(Y) :- q(Y, Y).
q(A, B) :- r(A), B = f(A).
r(42).
```

Here the relation q calls the function f , which in turn relies on the negation of the relation p (since the behavior of f is conditioned on the contents of p). As p is defined in terms of q , this leads to a

Function and expression well formedness
 $\Delta; \Phi \vdash F$
 $\Delta; \Phi; \Gamma \vdash e : \tau$

$$\begin{array}{c}
 \frac{f : \forall \vec{\alpha}_j, \vec{\tau}_i \rightarrow \tau \in \Phi \quad \vec{\alpha}_j, \vec{X}_i : \vec{\tau}_i \vdash e : \tau}{\vdash \text{fun } f(\vec{X}_i : \vec{\tau}_i) : \tau = e} \quad F\text{-WF} \\
 \\
 \frac{\vdash \Gamma \quad \Gamma(X) = \tau}{\Gamma \vdash X : \tau} \quad e\text{-VAR} \quad \frac{}{\Gamma \vdash k : \text{typeof}(k)} \quad e\text{-CONST} \quad \frac{\Gamma \vdash e_1 : \tau_1 \quad \Gamma, X : \tau_1 \vdash e_2 : \tau_2}{\Gamma \vdash \text{let } X = e_1 \text{ in } e_2 : \tau_2} \quad e\text{-LET} \\
 \\
 \frac{\Delta(D) = \forall \vec{\alpha}_j, \{\dots, c : \vec{\tau}_i, \dots\} \quad \Gamma \vdash_{\text{exp}} \tau'_j \quad \Gamma \vdash e_i : \tau_i[\tau'_j/\alpha_j]}{\Gamma \vdash c(\vec{e}_i) : D \vec{\tau}'_j} \quad e\text{-CTOR} \quad \frac{\Gamma \vdash \phi : \tau}{\Gamma \vdash \backslash \phi \backslash : \tau} \quad e\text{-QUOTE} \\
 \\
 \frac{p \subseteq \vec{\tau}_i \in \Phi \quad \Gamma \vdash e_i : \tau_i}{\Gamma \vdash p(\vec{e}_i) : \text{bool}} \quad e\text{-REL} \quad \frac{f : \forall \vec{\alpha}_j, \vec{\tau}_i \rightarrow \tau \in \Phi \quad \Gamma \vdash_{\text{exp}} \tau'_j \quad \Gamma \vdash e_i : \tau_i[\tau'_j/\alpha_j]}{\Gamma \vdash f(\vec{e}_i) : \tau[\tau'_j/\alpha_j]} \quad e\text{-FUN} \\
 \\
 \frac{\text{typeof}(\otimes) = \forall \vec{\alpha}_j, \vec{\tau}_i \rightarrow \tau \quad \Gamma \vdash_{\text{exp}} \tau'_j \quad \Gamma \vdash e_i : \tau_i[\tau'_j/\alpha_j]}{\Gamma \vdash \otimes(\vec{e}_i) : \tau} \quad e\text{-OP} \\
 \\
 \frac{\Gamma \vdash e_1 : \text{bool} \quad \Gamma \vdash e_2 : \tau \quad \Gamma \vdash e_3 : \tau}{\Gamma \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 : \tau} \quad e\text{-IF} \\
 \\
 \frac{\Gamma \vdash e : D \vec{\tau}_j \quad \Delta(D) = \forall \vec{\alpha}_j, \{\dots, c_i : \vec{\tau}_k, \dots\} \quad \Gamma, \vec{X}_k : \vec{\tau}_k[\tau_j/\alpha_j] \vdash e_i : \tau}{\Gamma \vdash \text{match } e \text{ with } c_i(\vec{X}_k) \rightarrow e_i : \tau} \quad e\text{-MATCH}
 \end{array}$$

Fig. 17. Typing rules: expressions; implicit parameters are in gray

circularity we want to avoid: It is possible to derive $q(42, 42)$, but this derivation nonsensically relies on $q(42, 42)$ being false (since it requires that $p(42)$ is false).

While Formulog's expressions *compute* values, the Formulog's formulas *construct* ASTs, to be shipped off to an SMT solver. Our formal account here uniformly uses SMT constructors to model the SMT syntax, but our implementation offers special-purpose syntax. For example, we write $c_{\text{var}}^{\text{SMT}}[x, \text{bv}[32]]$ in our formalism to name a 32-bit integer variable x , while in our implementation one might write $\#x[\text{bv}[32]]$.

Every ϕ -... typing rule generates a value with an SMT type, i.e., either $t \text{ sym}$ or $t \text{ smt}$ for SMT types t , i.e., $\Gamma \vdash_{\text{smt}} t$ (Lemma D.9). The c -SMT-* rules yield $t \text{ sym}$ and $t \text{ smt}$. SMT variables $c_{\text{var}}^{\text{SMT}}[x, t]$ are written in lowercase to emphasize their distinction from expression variables X ; these SMT variables will be used as names in the formulas sent to the SMT solver. We keep track of which terms are SMT variables $c_{\text{var}}^{\text{SMT}}[x, t]$ of type $t \text{ sym}$ (generated by c -SMT-VAR) and which are plain SMT formulas of type $t \text{ smt}$ (all other rules). We treat $t \text{ sym}$ as a subtype of $t \text{ smt}$ (ϕ -PROMOTE).

The $\backslash \phi \backslash$ operator is an expression term that introduces quoted SMT formulas represented as special, SMT constructors of the form c_{\dots}^{SMT} (described below); the $\text{,}e$ operator is the corresponding 'unquote' operator that introduces an expression (ϕ -UNQUOTE).

While unquoting generally suffices for embedding the results of expressions in formulas, we treat constructors specially so that we can mix concrete and symbolic (i.e., SMT) arguments in a single data type constructor (c -SMT-CTOR): we assign them types that are fully SMT-ized via the

SMT constructors and formula well formedness

$$\boxed{\Delta; \Phi; \Gamma \vdash c_{\dots}^{\text{SMT}} : \vec{\tau}_i \rightarrow \tau}$$

$$\boxed{\Delta; \Phi; \Gamma \vdash \phi : \tau}$$

$$\frac{\Gamma \vdash_{\text{smt}} t}{\Gamma \vdash c_{\text{var}}^{\text{SMT}}[x, t] : \cdot \rightarrow t \text{ sym}} \text{ c-SMT-VAR}$$

$$\frac{}{\Gamma \vdash c_{\text{const}}^{\text{SMT}}[k] : \cdot \rightarrow \text{typeof}(k) \text{ smt}} \text{ c-SMT-CONST}$$

$$\frac{\Gamma \vdash_{\text{smt}} t_1 \quad \Gamma \vdash_{\text{smt}} t_2}{\Gamma \vdash c_{\text{let}}^{\text{SMT}} : t_1 \text{ sym} \times t_1 \text{ smt} \times t_2 \text{ smt} \rightarrow t_2 \text{ smt}} \text{ c-SMT-LET}$$

$$\frac{\Delta(D) = \forall \vec{\alpha}_j, \{\dots, c : \vec{\tau}_i, \dots\} \quad \Gamma \vdash_{\text{smt}} \tau_i[t'_j/\alpha_j] \quad \Gamma \vdash_{\text{smt}} t'_j}{\Gamma \vdash c_{\text{ctor}}^{\text{SMT}}[c] : \overrightarrow{\text{toSMT}(\tau_i[t'_j/\alpha_j])} \rightarrow (D \vec{t}'_j) \text{ smt}} \text{ c-SMT-CTOR}$$

$$\frac{\Gamma \vdash_{\text{smt}} t_1}{\Gamma \vdash c_{\text{forall}}^{\text{SMT}} : t_1 \text{ sym} \times \text{bool smt} \rightarrow \text{bool smt}} \text{ c-SMT-FORALL}$$

$$\frac{uf : \vec{t}_i \rightarrow t \in \Phi}{\Gamma \vdash c_{\text{uf}}^{\text{SMT}}[uf] : \vec{t}_i \text{ smt} \rightarrow t \text{ smt}} \text{ c-SMT-UFUN}$$

$$\frac{\Gamma \vdash \phi : t \text{ sym}}{\Gamma \vdash \phi : t \text{ smt}} \text{ } \phi\text{-PROMOTE}$$

$$\frac{\Gamma \vdash e : \tau \quad \Gamma \vdash_{\text{smt}} \tau}{\Gamma \vdash ,e : \text{toSMT}(\tau)} \text{ } \phi\text{-UNQUOTE}$$

$$\frac{\Gamma \vdash c_c^{\text{SMT}} : \vec{\tau}_i \rightarrow \tau \quad \Gamma \vdash \phi_i : \tau'_i}{\Gamma \vdash c_c^{\text{SMT}}(\vec{\phi}_i) : \tau} \text{ } \phi\text{-CTOR}$$

Conversion to SMT types

$$\boxed{\text{erase}(\tau) = t}$$

$$\boxed{\text{toSMT}(\tau) = \tau}$$

$$\text{erase}(B) = B$$

$$\text{erase}(D \vec{\tau}_i) = D \overrightarrow{\text{erase}(\tau_i)}$$

$$\text{erase}(t \text{ smt}) = \text{erase}(t)$$

$$\text{erase}(t \text{ sym}) = \text{erase}(t)$$

$$\text{toSMT}(t) = \text{erase}(t) \text{ smt}$$

$$\text{toSMT}(t \text{ smt}) = \text{erase}(t) \text{ smt}$$

$$\text{toSMT}(t \text{ sym}) = \text{erase}(t) \text{ sym}$$

Fig. 18. Typing rules: SMT constructors and formulas; conversion to SMT types

toSMT function, but unquoting allows for easy mixing of values of SMT-types t as if they were of type $t \text{ smt}$. The toSMT metafunction alters the type of e to make sure it is SMT representable; toSMT relies on an erase function to avoid nesting $\dots \text{ smt}$ and $\dots \text{ sym}$ type constructors. One can only run these functions on SMT types. For example, we can write terms like the following in concrete syntax:

```
let H = 5 in `cons(H, #l[bv[32] list])`
```

which desugars to the SMT constructors:

```
let H = 5 in `c_{ctor}^{\text{SMT}}[\text{cons}](,H, c_{\text{var}}^{\text{SMT}}[l, \text{bv}[32] \text{ list}]())`
```

Note that H is an expression variable and l is an SMT variable; the type conversion in $\phi\text{-UNQUOTE}$ lets us mix them in the same list of 32-bit numbers. The $t \text{ sym}$ type is used in $\phi\text{-LET}$ and $\phi\text{-FORALL}$, which construct SMT formulae that use binders. The only way to get a value of type $t \text{ sym}$ is either with $c\text{-SMT-VAR}/\phi\text{-CTOR}$ or with $\phi\text{-UNQUOTE}$, as in `let X = `c_{var}^{\text{SMT}}[x, bv[32]]` in `X``.

Uninterpreted functions must be applied to appropriate SMT types (ϕ -UFUN); recall that Φ -UFUN ensures that each uninterpreted function's types are SMT representable.

Finally, there are a suite of SMT constructors of the form c_{\dots}^{SMT} . Each of these special c_{\dots}^{SMT} constructors is treated as an ordinary constructor by the operational semantics, even though the constructors don't appear in Δ . Rather than making SMT terms opaque, we model them with constructors to allow for matching on generated formulae in P -EQSMT-BF and P -EQ-FB. The types of the SMT constructors c_{\dots}^{SMT} are given in Figure 18. It is a crucial invariant that all of these types be SMT types: we would not want to treat an SMT variable $c_{\text{var}}^{\text{SMT}}[x, \text{bool}]$ as though it were an *actual* bool! Reusing the c -SMT-... rules is convenient—we need only state the types of these constructors once and we get precise types in our premises. Several SMT constructors take special arguments in square brackets: $c_{\text{var}}^{\text{SMT}}[x, t]$ is a 0-ary SMT constructor, while $c_{\text{var}}^{\text{SMT}}$ itself is a family of SMT constructors for given variable names x and pre-types t ; similarly, $c_{\text{const}}^{\text{SMT}}[k]$ is a 0-ary SMT constructor, while $c_{\text{const}}^{\text{SMT}}$ itself is a family of SMT constructors for given constants k . The embedding of data type constructors $c_{\text{ctor}}^{\text{SMT}}[c]$ is similarly parameterized on a constructor name c , and the embedding of uninterpreted functions $c_{\text{uf}}^{\text{SMT}}[uf]$ takes an uninterpreted function as a parameter. Separating these parameters from the interesting, ϕ -shaped subparts of each SMT constructor lets us reuse the c -SMT-... rules when typing premises that might bind to subparts of an SMT formula (P -EQSMT-... , Figure 16).

Our formal model elides some of the detail of our SMT encoding, such as constructors for SMT operations like bit vector manipulation or equality. These operations are all encoded as more SMT-specific constructors, i.e., $c_{\text{ctor}}^{\text{SMT}}[\text{bv}_{32_add}](c_{\text{var}}^{\text{SMT}}[x, \text{bv}[32]], c_{\text{const}}^{\text{SMT}}[1])$ represents result of adding the 32-bit vector x and 1. There are some subtle issues around polymorphism and determinacy. We give a monomorphic interpretation of SMT here, though our SMT constructors can work with our polymorphic data types. Our implementation treats equality and other polymorphic SMT operations specially, where each use of a polymorphic operator must be fully instantiated. In practice our implementation can usually infer the instantiation; users must annotate in those places we cannot infer.

C OPERATIONAL SEMANTICS

Formulog's operational semantics operates over *worlds* \mathcal{W} and substitutions θ (Figure 19); the semantics is a mix of small-step rules modeling a single application of a Datalog rule (Figure 20), which depend on a small-step rules explaining how premises unify (Figures 21 and 22); the premise semantics in turn depends on a semantics of expressions (Figures 23 and 24) and formulas (Figure 25).

Our worlds \mathcal{W} are (subsets of) Herbrand models. Our small-step semantics iteratively builds up a world that is in fact a Herbrand model of the original relations in the program. We could have modeled our semi-naive evaluation model for Formulog in more detail, showing that all programs generate a world \mathcal{W} that is a well typed Herbrand model of the user's program (possibly taking infinite time to do so). Doing so wouldn't add anything materially interesting to our formulation.

Throughout, the type system's goal is prevent a program yielding \perp , the bottom "wrong" value. Such a value denotes a serious, unrecoverable error, such as using a relation with the wrong arity or conditioning on a non-boolean. It is important to distinguish bad, \perp -yielding programs from those that simply fail to step. The goal of Datalog evaluation is to reach a fixed point, i.e., to be unable to step! Finally, as is common, we assume that built-in operations do not yield \perp , i.e., they are total. While we could in principle design a type system for Formulog that avoids, say, division by zero, we are more interested in making the hard parts easy (generating well typed SMT formulas) rather than making the easy parts foolproof (statically protecting partial functions).

Namespaces

World	\mathcal{W}	\in	$\text{PredVar} \rightarrow \mathcal{P}(\text{Val} \times \dots \times \text{Val})$
World or error	\mathcal{W}_\perp	\in	World + Error
Substitution	θ	\in	$\text{Var} \rightarrow \text{Val}$
Substitution or error	θ_\perp	\in	Substitution + Error

Values

Results	v_\perp	$::=$	$v \mid \perp$
Values	$v \in \text{Val}$	$::=$	$k \mid c(\vec{v}_i)$

Unifiable term	$u \in \text{UTerm}$	$::=$	$X \mid k \mid c(\vec{u}_i)$
----------------	----------------------	-------	------------------------------

Substitution and world well formedness

	$\Delta; \Phi; \Gamma \models \theta$	$\Delta; \Phi \models \mathcal{W}$
	\Leftrightarrow	\Leftrightarrow
$\forall X \in \text{dom}(\Gamma) \left\{ \begin{array}{l} (1) X \in \text{dom}(\theta) \\ (2) \cdot \vdash \theta(X) : \Gamma(X) \end{array} \right.$		$\forall p \subseteq \vec{\tau}_i \in \Phi \left\{ \begin{array}{l} (1) p \in \text{dom}(\mathcal{W}) \\ (2) \vec{v}_j \in \mathcal{W}(p) \Rightarrow i = j \\ (3) \forall \vec{v}_i \in \mathcal{W}(p), \Delta; \Phi; \cdot \vdash v_i : \tau_i \end{array} \right.$

Fig. 19. Definitions for semantics

Clause semantics

	$\vec{F}; \mathcal{W} \vdash H \rightarrow \mathcal{W}_\perp$
$\frac{\cdot \vdash P_0 \rightarrow \theta_1 \quad \dots \quad \theta_i \vdash P_i \rightarrow \theta_{i+1} \quad \dots \quad \theta_n \vdash P_n \rightarrow \theta}{\vec{F}; \mathcal{W} \vdash p(\vec{X}_j) :- \vec{P}_i \rightarrow \mathcal{W}[p \mapsto \mathcal{W}(p) \cup \{\theta(\vec{X}_j)\}]}$	CLAUSE
$\frac{\cdot \vdash P_0 \rightarrow \theta_1 \quad \dots \quad \theta_i \vdash P_j \rightarrow \perp}{\vec{F}; \mathcal{W} \vdash p(\vec{X}_j) :- \vec{P}_i \rightarrow \perp}$	CLAUSE-E1
$\frac{\cdot \vdash P_0 \rightarrow \theta_1 \quad \dots \quad \theta_i \vdash P_i \rightarrow \theta_{i+1} \quad \dots \quad \theta_n \vdash P_n \rightarrow \theta \quad \vec{X}_j \not\subseteq \text{dom}(\theta)}{\vec{F}; \mathcal{W} \vdash p(\vec{X}_j) :- \vec{P}_i \rightarrow \perp}$	CLAUSE-E2

Fig. 20. Clause semantics

Rules of the form \dots -En denote \perp -yielding rules. Each such rule characterizes a form of wrongness avoided by our static type system. We write v_\perp to denote the disjoint sum of values v and the wrong value \perp .

During correct execution, CLAUSE takes a Horn clause $p(\vec{X}_j) :- \vec{P}_i$, executes each premise P_i from left to right, yielding a final substitution for the variables \vec{X}_j in the head of the rule. There are two possible failing rules. CLAUSE-E1 simply propagates the first error from a premise; CLAUSE-E2 fails because not every X_j in the head of the rule is bound by the end. Since Formulog enforces the range restriction (H-Clause), CLAUSE-E2 can never apply in a well typed program. Finally, the CLAUSE* operational rules and the H-CLAUSE typing rule both use the fixed, given order of premises

Premise semantics

$$\vec{F}; \mathcal{W}; \theta \vdash P \rightarrow \theta_{\perp}$$

$$\begin{array}{c}
\frac{\vec{v} \in \mathcal{W}(p) \quad \theta \vdash \vec{X} \sim \vec{v} : \theta'_{\perp}}{\mathcal{W}; \theta \vdash p(\vec{X}) \rightarrow \theta'_{\perp}} \text{PosATOM} \qquad \frac{\theta(\vec{X}) = \vec{v} \quad \vec{v} \notin \mathcal{W}(p)}{\mathcal{W}; \theta \vdash !p(\vec{X}) \rightarrow \theta} \text{NEGATOM} \\
\frac{\theta \vdash Y \sim c(\vec{X}) : \theta'_{\perp}}{\mathcal{W}; \theta \vdash Y = c(\vec{X}) \rightarrow \theta'_{\perp}} \text{EqCTOR} \qquad \frac{\theta \vdash Y \sim c_c^{\text{SMT}}(\vec{X}) : \theta'_{\perp}}{\mathcal{W}; \theta \vdash Y = `c_c^{\text{SMT}}(\vec{X})` \rightarrow \theta'_{\perp}} \text{EqSMT} \\
\frac{e \text{ is not a constructor} \quad \mathcal{W}; \theta \vdash e \Downarrow_e v \quad \theta \vdash Y \sim v : \theta'_{\perp}}{\mathcal{W}; \theta \vdash Y = e \rightarrow \theta'_{\perp}} \text{EqEXPR} \\
\frac{\vec{X} \not\subseteq \text{dom}(\theta)}{\mathcal{W}; \theta \vdash !p(\vec{X}) \rightarrow \perp} \text{NEGATOM-E} \qquad \frac{e \text{ is not a constructor} \quad \mathcal{W}; \theta \vdash e \Downarrow_e \perp}{\mathcal{W}; \theta \vdash Y = e \rightarrow \perp} \text{EqEXPR-E} \\
\frac{\mathcal{W}; \theta \vdash e \Downarrow_e v \quad \theta(Y) \neq v}{\mathcal{W}; \theta \vdash !(Y = e) \rightarrow \theta} \text{NEGEXPR} \qquad \frac{\mathcal{W}; \theta \vdash e \Downarrow_e \perp}{\mathcal{W}; \theta \vdash !(Y = e) \rightarrow \perp} \text{NEGEXPR-E1} \qquad \frac{Y \notin \text{dom}(\theta)}{\mathcal{W}; \theta \vdash !(Y = e) \rightarrow \perp} \text{NEGEXPR-E2}
\end{array}$$

Fig. 21. Premise semantics

for checking. Different orderings induce different binding orders, some of which may succeed and some of which may not.

The premise semantics (Figure 21) uses unification (Figure 22) to match and bind variables. Positive atoms $p(\vec{X}_i)$ try to unify their arguments with a tuple for p drawn from the world \mathcal{W} (PosATOM). Negative atoms $p(\vec{X}_i)$ require that all of their arguments X_i are already bound (NEGATOM); failing to find such bound terms yields an error (NEGATOM-E). Rules for equations also use unification, whether for a constructor over variables (EqCTOR) or an expression (EqEXPR). The latter can fail if evaluation fails (EqEXPR-E). Before discussing term evaluation, we give rules for unification.

Unification is split into two levels. *Unification* proper takes a pair of unifiable terms—values with variables in them—and tries to yield a substitution. *Value unification* takes a unifiable term and a value and tries to yield a substitution. The unification rules are of the form $uu-\dots$. These rules analyze the two unifiable terms to find which side is completely bound—i.e., applying θ can completely fill in the variables—and so can be passed to value unification as a value. The B and F in these rules stand for Bound and Free. The only error in unification is in uu -FF, when neither unifiable term is bound to a value. We write a case for when both unifiable terms are bound (uu -BB) and require that they are directly equal—but it would also work to drop this rule and rely on value unification to identify the equality.

Value unification rules are of the form $uv-\dots$. The rules here lookup variables in the unifiable term and either check that the binding conforms to the given value (uv -EQ-VAR, cf. $X\tau$ -CHECK) or binds the value (uv -BIND-VAR, cf. $X\tau$ -BIND). The remaining value unification rules match the structure of the unifiable term to the structure of the value (uv -CONSTANT, uv -CTOR) or fold value unification along a vector ($\vec{u}\vec{v}$ -ALL). Value unification never produces \perp . It isn't an error when two values fail to unify, since one might have to search through many tuples for a relation in \mathcal{W} to find a one that matches, say, a given constructor.

Value unification

$$\begin{array}{c}
\boxed{\theta \vdash u \sim v \triangleright \theta} \quad \boxed{\theta \vdash \vec{u} \sim \vec{v} \triangleright \theta} \\
\\
\frac{\theta(X) = v}{\theta \vdash X \sim v \triangleright \theta} \quad uv\text{-EQ-VAR} \qquad \frac{X \notin \text{dom}(\theta)}{\theta \vdash X \sim v \triangleright \theta[X \mapsto v]} \quad uv\text{-BIND-VAR} \\
\\
\frac{}{\theta \vdash k \sim k \triangleright \theta} \quad uv\text{-CONSTANT} \qquad \frac{\theta \vdash \vec{u}_i \sim \vec{v}_i \triangleright \theta'}{\theta \vdash c(\vec{u}_i) \sim c(\vec{v}_i) \triangleright \theta'} \quad uv\text{-CTOR} \\
\\
\frac{\theta \vdash u_0 \sim v_0 \triangleright \theta_1 \quad \dots \quad \theta_1 \vdash u_i \sim v_i \triangleright \theta_i \quad \dots \quad \theta_n \vdash u_n \sim v_n \triangleright \theta'}{\theta \vdash \vec{u}_i \sim \vec{v}_i \triangleright \theta} \quad \vec{u}\vec{v}\text{-ALL}
\end{array}$$

Unification

$$\begin{array}{c}
\boxed{\theta \vdash u \sim u : \theta_{\perp}} \quad \boxed{\theta \vdash \vec{u} \sim \vec{u} : \theta_{\perp}} \\
\\
\frac{\theta(u_1) = v_1 \quad \theta(u_2) = v_2 \quad v_1 = v_2}{\theta \vdash u_1 \sim u_2 : \theta} \quad uu\text{-BB} \qquad \frac{\nexists v_1, \theta(u_1) = v_1 \quad \theta(u_2) = v_2 \quad \theta \vdash u_1 \sim v_2 \triangleright \theta'}{\theta \vdash u_1 \sim u_2 : \theta'} \quad uu\text{-FB} \\
\\
\frac{\theta(u_1) = v_1 \quad \nexists v_2, \theta(u_2) = v_2 \quad \theta \vdash u_2 \sim v_1 \triangleright \theta'}{\theta \vdash u_1 \sim u_2 : \theta'} \quad uu\text{-BF} \qquad \frac{\nexists v_1, \theta(u_1) = v_1 \quad \nexists v_2, \theta(u_2) = v_2}{\theta \vdash u_1 \sim u_2 : \perp} \quad uu\text{-FF} \\
\\
\frac{\theta \vdash u_i \sim u'_i \triangleright \theta_i}{\theta \vdash \vec{u}_i \sim \vec{u}'_i \triangleright \theta\vec{\theta}_i} \quad \vec{u}\vec{u}\text{-ALL} \qquad \dots \quad \frac{\theta \vdash u \sim u' \triangleright \perp \quad \dots}{\theta \vdash \vec{u}_i \sim \vec{u}'_i \triangleright \perp} \quad \vec{u}\vec{u}\text{-ALL-E} \\
\\
\theta(c(\vec{u}_i)) = c(\overrightarrow{\theta(u_i)})
\end{array}$$

Fig. 22. Unification

The expression semantics is an entirely conventional big-step semantics using explicit substitutions. The operational rules implicitly take the function definitions \vec{F} for use in applications ($\Downarrow_e\text{-FUN}$).

There are a variety of wrong behaviors prevented by our type system, mostly concerning mismatches between values and elimination forms: unbound variables ($\Downarrow_e\text{-VAR-E}$); mistyped arguments to built-in operations ($\Downarrow_e\text{-OP-E2}$); function, relation and constructor arity errors ($\Downarrow_e\text{-FUN-E2}$, $\Downarrow_e\text{-REL-E2}$, $\Downarrow_e\text{-MATCH-E4}$); non-existent functions and relations ($\Downarrow_e\text{-FUN-E3}$, $\Downarrow_e\text{-REL-E3}$); conditionals on inappropriate values ($\Downarrow_e\text{-MATCH-E2}$, $\Downarrow_e\text{-ITE-E2}$); and ill formed constructor names ($\Downarrow_e\text{-MATCH-E3}$); The remaining rules propagate errors ($\Downarrow_e\text{-LET-E}$, $\Downarrow_e\text{-OP-E1}$, $\Downarrow_e\text{-FUN-E1}$, $\Downarrow_e\text{-REL-E1}$, $\Downarrow_e\text{-MATCH-E1}$, $\Downarrow_e\text{-ITE-E1}$). As mentioned in the early discussion of our semantics in this section, $\Downarrow_e\text{-OP-E2}$ is *not* about division by zero (a form of going wrong our type system *doesn't* prevent), but about mis-application of built-in functions, e.g., taking the boolean negation of a number.

The operational semantics on formulas is simple: the rules generate ASTs for the SMT solver using the c^{SMT} constructors: constants ($c_{\text{const}}^{\text{SMT}}$), SMT variables ($c_{\text{var}}^{\text{SMT}}$), SMT data types ($c_{\text{ctor}}^{\text{SMT}}$), let bindings ($c_{\text{let}}^{\text{SMT}}$), quantification ($c_{\text{forall}}^{\text{SMT}}$), and uninterpreted function application ($c_{\text{uf}}^{\text{SMT}}$).

When unquoting values resulting from evaluating expressions, we use the toSMT function to translate expression values into the SMT's AST. The toSMT function is an identity on SMT ASTs, but it explicitly tags the constants and Formolog-defined constructors using $c_{\text{const}}^{\text{SMT}}$ and $c_{\text{ctor}}^{\text{SMT}}$.

Expression semantics

$$\begin{array}{c}
\boxed{\vec{F}; \mathcal{W}; \theta \vdash e \Downarrow_e v_\perp} \quad \boxed{\vec{F}; \mathcal{W}; \theta \vdash \vec{e} \Downarrow_{\vec{e}} \vec{v}_\perp} \\
\\
\frac{}{\mathcal{W}; \theta \vdash \cdot \Downarrow_{\vec{e}} \cdot} \Downarrow_{\vec{e}}\text{-EMPTY} \qquad \frac{\mathcal{W}; \theta \vdash e \Downarrow_e v \quad \mathcal{W}; \theta \vdash \vec{e} \Downarrow_{\vec{e}} \vec{v}}{\mathcal{W}; \theta \vdash e, \vec{e} \Downarrow_{\vec{e}} v, \vec{v}} \Downarrow_{\vec{e}}\text{-ALL} \\
\\
\frac{\mathcal{W}; \theta \vdash \vec{e}_i \Downarrow_{\vec{e}} \vec{v}_i}{\mathcal{W}; \theta \vdash c(\vec{e}_i) \Downarrow_{\vec{e}} c(\vec{v}_i)} \Downarrow_{\vec{e}}\text{-CTOR} \qquad \frac{\mathcal{W}; \theta \vdash \vec{e}_i \Downarrow_{\vec{e}} \perp}{\mathcal{W}; \theta \vdash c(\vec{e}_i) \Downarrow_{\vec{e}} \perp} \Downarrow_{\vec{e}}\text{-CTOR-E} \\
\\
\frac{}{\mathcal{W}; \theta \vdash k \Downarrow_e k} \Downarrow_e\text{-CONST} \qquad \frac{\theta(X) = v}{\mathcal{W}; \theta \vdash X \Downarrow_e v} \Downarrow_e\text{-VAR} \qquad \frac{\mathcal{W}; \theta \vdash \phi \Downarrow_\phi v_\perp}{\mathcal{W}; \theta \vdash \text{'}\phi\text{'} \Downarrow_e v_\perp} \Downarrow_e\text{-QUOTE} \\
\\
\frac{\mathcal{W}; \theta \vdash \vec{e} \Downarrow_{\vec{e}} \vec{v} \quad \llbracket \otimes \rrbracket(\vec{v}) = v}{\mathcal{W}; \theta \vdash \otimes(\vec{e}) \Downarrow_e v} \Downarrow_e\text{-OP} \\
\\
\frac{\text{fun } f(\vec{X}_i : \vec{\tau}_i) : \tau = e \in \vec{F} \quad \mathcal{W}; \theta \vdash \vec{e}_i \Downarrow_{\vec{e}} \vec{v}_i \quad \mathcal{W}; \theta[\vec{X}_i \mapsto \vec{v}_i] \vdash e \Downarrow_e v_\perp}{\mathcal{W}; \theta \vdash f(\vec{e}_i) \Downarrow_e v_\perp} \Downarrow_e\text{-FUN} \\
\\
\frac{\mathcal{W}; \theta \vdash \vec{e}_i \Downarrow_{\vec{e}} \vec{v}_i \quad \vec{v}_i \in \mathcal{W}(p)}{\mathcal{W}; \theta \vdash p(\vec{e}_i) \Downarrow_e \text{true}} \Downarrow_e\text{-REL-T} \qquad \frac{\mathcal{W}; \theta \vdash \vec{e}_i \Downarrow_{\vec{e}} \vec{v}_i \quad \vec{v}_i \notin \mathcal{W}(p)}{\mathcal{W}; \theta \vdash p(\vec{e}_i) \Downarrow_e \text{false}} \Downarrow_e\text{-REL-F} \\
\\
\frac{\mathcal{W}; \theta \vdash e_1 \Downarrow_e v_1 \quad \mathcal{W}; \theta[X \mapsto v_1] \vdash e_2 \Downarrow_e v_\perp}{\mathcal{W}; \theta \vdash \text{let } X = e_1 \text{ in } e_2 \Downarrow_e v_\perp} \Downarrow_e\text{-LET} \\
\\
\frac{\mathcal{W}; \theta \vdash e \Downarrow_e c(\vec{v}_i) \quad \mathcal{W}; \theta[\vec{X}_i \mapsto \vec{v}_i] \vdash e \Downarrow_e v_\perp}{\mathcal{W}; \theta \vdash \text{match } e \text{ with } \dots c(\vec{X}_i) \rightarrow e \dots \Downarrow_e v_\perp} \Downarrow_e\text{-MATCH} \\
\\
\frac{\mathcal{W}; \theta \vdash e_1 \Downarrow_e \text{true} \quad \mathcal{W}; \theta \vdash e_2 \Downarrow_e v_\perp}{\mathcal{W}; \theta \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \Downarrow_e v_\perp} \Downarrow_e\text{-ITET} \qquad \frac{\mathcal{W}; \theta \vdash e_1 \Downarrow_e \text{false} \quad \mathcal{W}; \theta \vdash e_3 \Downarrow_e v_\perp}{\mathcal{W}; \theta \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \Downarrow_e v_\perp} \Downarrow_e\text{-ITEF}
\end{array}$$

Fig. 23. Expression semantics

D METATHEORY

We break the metatheory into two parts: lemmas characterizing the SMT conversion (Section D.1) and lemmas showing type safety (Section E). The SMT lemmas culminate in two proofs: first, regularity (Lemma D.9) guarantees that (a) every type or context generated by the operational semantics is well formed and (b) that formula evaluation generates well typed SMT ASTs; second, we show that SMT conversion of values agrees with SMT conversion of types (Lemma D.10). Type safety culminates in theorems showing that premises don't yield \perp and generate well typed substitutions (Lemma E.4) and so Horn clauses (a) never yield \perp (Theorem E.6) and (b) take well typed worlds to well typed worlds (Theorem E.7).

D.1 SMT conversion

We show a variety of properties of the erasure and SMT conversion functions: `smt` is a sub-kind of `exp` (Lemma D.1); erasures and SMT conversion yield well formed types from SMT types

Expression semantics (continued)

$$\boxed{\vec{F}; \mathcal{W}; \theta \vdash e \Downarrow_e v \perp}$$

$$\boxed{\vec{F}; \mathcal{W}; \theta \vdash \vec{e} \Downarrow_{\vec{e}} \vec{v} \perp}$$

$$\begin{array}{c}
\frac{\mathcal{W}; \theta \vdash \vec{e}_i \Downarrow_{\vec{e}} \vec{v}_i \quad \mathcal{W}; \theta \vdash e \Downarrow_e \perp}{\mathcal{W}; \theta \vdash \vec{e}_i, e, \vec{e}_j \Downarrow_{\vec{e}} \perp} \Downarrow_{\vec{e}}\text{-ALL-E} \\
\\
\frac{X \notin \text{dom}(\theta)}{\mathcal{W}; \theta \vdash X \Downarrow_e \perp} \Downarrow_e\text{-VAR-E} \qquad \frac{\mathcal{W}; \theta \vdash e_1 \Downarrow_e \perp}{\mathcal{W}; \theta \vdash \text{let } X = e_1 \text{ in } e_2 \Downarrow_e \perp} \Downarrow_e\text{-LET-E} \\
\\
\frac{\mathcal{W}; \theta \vdash \vec{e} \Downarrow_{\vec{e}} \perp}{\mathcal{W}; \theta \vdash \otimes(\vec{e}) \Downarrow_e \perp} \Downarrow_e\text{-OP-E1} \qquad \frac{\mathcal{W}; \theta \vdash \vec{e} \Downarrow_{\vec{e}} \vec{v} \quad \vec{v} \notin \text{dom}(\llbracket \otimes \rrbracket)}{\mathcal{W}; \theta \vdash \otimes(\vec{e}) \Downarrow_e \perp} \Downarrow_e\text{-OP-E2} \\
\\
\frac{\mathcal{W}; \theta \vdash \vec{e}_i \Downarrow_{\vec{e}} \perp}{\mathcal{W}; \theta \vdash f(\vec{e}_i) \Downarrow_e \perp} \Downarrow_e\text{-FUN-E1} \qquad \frac{\text{fun } f(\vec{X}_i : \vec{\tau}_i) : \tau = e \in \vec{F} \quad i \neq j}{\mathcal{W}; \theta \vdash f(\vec{e}_j) \Downarrow_e \perp} \Downarrow_e\text{-FUN-E2} \qquad \frac{f \notin \vec{F}}{\mathcal{W}; \theta \vdash f(\vec{e}_j) \Downarrow_e \perp} \Downarrow_e\text{-FUN-E3} \\
\\
\frac{\mathcal{W}; \theta \vdash \vec{e}_i \Downarrow_{\vec{e}} \perp}{\mathcal{W}; \theta \vdash p(\vec{e}_i) \Downarrow_e \perp} \Downarrow_e\text{-REL-E1} \qquad \frac{\mathcal{W}(p) \subseteq \mathcal{P}(\overrightarrow{\text{Val}}_i) \quad i \neq j}{\mathcal{W}; \theta \vdash p(\vec{e}_j) \Downarrow_e \perp} \Downarrow_e\text{-REL-E2} \qquad \frac{p \notin \text{dom}(\mathcal{W})}{\mathcal{W}; \theta \vdash p(\vec{e}_j) \Downarrow_e \perp} \Downarrow_e\text{-REL-E3} \\
\\
\frac{\mathcal{W}; \theta \vdash e \Downarrow_e \perp}{\mathcal{W}; \theta \vdash \text{match } e \text{ with } c_i(\vec{X}_j) \rightarrow e_i \Downarrow_e \perp} \Downarrow_e\text{-MATCH-E1} \qquad \frac{\mathcal{W}; \theta \vdash e \Downarrow_e v \quad v \neq c(\vec{v}')}{\mathcal{W}; \theta \vdash \text{match } e \text{ with } c_i(\vec{X}_j) \rightarrow e_i \Downarrow_e \perp} \Downarrow_e\text{-MATCH-E2} \\
\\
\frac{\mathcal{W}; \theta \vdash e \Downarrow_e c(\vec{v}_k) \quad c \notin \{\vec{c}_i\}}{\mathcal{W}; \theta \vdash \text{match } e \text{ with } c_i(\vec{X}_j) \rightarrow e_i \Downarrow_e \perp} \Downarrow_e\text{-MATCH-E3} \qquad \frac{\mathcal{W}; \theta \vdash e \Downarrow_e c(\vec{v}_k) \quad j \neq k}{\mathcal{W}; \theta \vdash \text{match } e \text{ with } \dots c(\vec{X}_j) \dots \rightarrow e_i \Downarrow_e \perp} \Downarrow_e\text{-MATCH-E4} \\
\\
\frac{\mathcal{W}; \theta \vdash e_1 \Downarrow_e \perp}{\mathcal{W}; \theta \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \Downarrow_e \perp} \Downarrow_e\text{-ITE-E1} \qquad \frac{\mathcal{W}; \theta \vdash e_1 \Downarrow_e v \quad v \notin \{\text{true}, \text{false}\}}{\mathcal{W}; \theta \vdash \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \Downarrow_e \perp} \Downarrow_e\text{-ITE-E2}
\end{array}$$

Fig. 24. Expression semantics (error rules)

(Lemmas D.2, D.3, D.4, and D.5); weakening and strengthening of typing contexts (Lemmas D.6 and D.7); type variable substitution (Lemma D.8)—we have no need of a value substitution lemma because our semantics uses environments; regularity (Lemma D.9); and, finally, that SMT conversion of values agrees with SMT conversion of types (Lemma D.10).

LEMMA D.1 (smt IS A SUBKIND OF exp). *If $\Gamma \vdash_{\text{smt}} \tau$ then $\Gamma \vdash_{\text{exp}} \tau$.*

PROOF. By induction on τ .

($\tau = B$) Immediate: τ -BASE allows any m .

($\tau = \alpha$) Contradictory—type variables are only well formed at exp.

($\tau = D \vec{\tau}_i$) By the IH on each τ_i and t -ADT.

Formula semantics

$$\begin{array}{c}
 \boxed{\vec{F}; \mathcal{W}; \theta \vdash \phi \Downarrow_{\vec{\phi}} v_{\perp}} \quad \boxed{\vec{F}; \mathcal{W}; \theta \vdash \vec{\phi} \Downarrow_{\vec{\phi}} \vec{v}_{\perp}} \\
 \\
 \frac{}{\mathcal{W}; \theta \vdash \cdot \Downarrow_{\vec{\phi}} \cdot} \Downarrow_{\vec{\phi}}\text{-EMPTY} \quad \frac{\mathcal{W}; \theta \vdash \phi \Downarrow_{\vec{\phi}} v \quad \mathcal{W}; \theta \vdash \vec{\phi}_i \Downarrow_{\vec{\phi}} \vec{v}_i}{\mathcal{W}; \theta \vdash \phi, \vec{\phi}_i \Downarrow_{\vec{\phi}} v, \vec{v}_i} \Downarrow_{\vec{\phi}}\text{-ALL} \\
 \\
 \frac{\mathcal{W}; \theta \vdash \vec{\phi}_i \Downarrow_{\vec{\phi}} \vec{v}_i \quad \mathcal{W}; \theta \vdash \phi \Downarrow_{\vec{\phi}} \perp}{\mathcal{W}; \theta \vdash \vec{\phi}_i, \phi, \vec{\phi}_j \Downarrow_{\vec{\phi}} \perp} \Downarrow_{\vec{\phi}}\text{-ALL-E} \\
 \\
 \frac{\mathcal{W}; \theta \vdash e \Downarrow_e v}{\mathcal{W}; \theta \vdash ,e \Downarrow_{\vec{\phi}} \text{toSMT}(v)} \Downarrow_{\vec{\phi}}\text{-UNQUOTE} \quad \frac{\mathcal{W}; \theta \vdash e \Downarrow_e \perp}{\mathcal{W}; \theta \vdash ,e \Downarrow_{\vec{\phi}} \perp} \Downarrow_{\vec{\phi}}\text{-UNQUOTE-E} \\
 \\
 \frac{\mathcal{W}; \theta \vdash \vec{\phi}_i \Downarrow_{\vec{\phi}} \vec{v}_i}{\mathcal{W}; \theta \vdash c_c^{\text{SMT}}(\vec{\phi}) \Downarrow_{\vec{\phi}} c_c^{\text{SMT}}(\vec{v})} \Downarrow_{\vec{\phi}}\text{-CTOR} \quad \frac{\mathcal{W}; \theta \vdash \vec{\phi}_i \Downarrow_{\vec{\phi}} \perp}{\mathcal{W}; \theta \vdash c_c^{\text{SMT}}(\vec{\phi}_i) \Downarrow_{\vec{\phi}} \perp} \Downarrow_{\vec{\phi}}\text{-CTOR-E} \\
 \\
 \frac{}{\mathcal{W}; \theta \vdash c_c^{\text{SMT}}(\vec{v}_i) \Downarrow_{\vec{\phi}} c_c^{\text{SMT}}(\vec{v}_i)} \Downarrow_{\vec{\phi}}\text{-SMT-VALUE}
 \end{array}$$

SMT conversion

$$\begin{array}{c}
 \boxed{\text{toSMT}(v) = v} \\
 \\
 \begin{array}{ll}
 \text{toSMT}(k) = c_{\text{const}}^{\text{SMT}}[k]() & \text{toSMT}(c_{\text{ctor}}^{\text{SMT}}[c](\vec{v}_i)) = c_{\text{ctor}}^{\text{SMT}}[c](\vec{v}_i) \\
 \text{toSMT}(c(\vec{v}_i)) = c_{\text{ctor}}^{\text{SMT}}[c](\text{toSMT}(\vec{v}_i)) & \text{toSMT}(c_{\text{let}}^{\text{SMT}}(v_1, v_2, v_3)) = c_{\text{let}}^{\text{SMT}}(v_1, v_2, v_3) \\
 \text{toSMT}(c_{\text{const}}^{\text{SMT}}[k]()) = c_{\text{const}}^{\text{SMT}}[k]() & \text{toSMT}(c_{\text{forall}}^{\text{SMT}}(v_1, v_2)) = c_{\text{forall}}^{\text{SMT}}(v_1, v_2) \\
 \text{toSMT}(c_{\text{var}}^{\text{SMT}}[x, t]()) = c_{\text{var}}^{\text{SMT}}[x, t]() & \text{toSMT}(c_{\text{uf}}^{\text{SMT}}[\text{uf}](\vec{v}_i)) = c_{\text{uf}}^{\text{SMT}}[\text{uf}](\vec{v}_i)
 \end{array}
 \end{array}$$

Fig. 25. Formula semantics

($\tau = t$ smt) Immediate: τ -SMT allows any m .

($\tau = t$ sym) Immediate: τ -SMT allows any m .

($\tau = \text{model}$) Contradictory—model is only well formed at exp. □

LEMMA D.2 (ERASURE IS WELL FORMED). *If $\Gamma \vdash_{\text{smt}} \tau$ then $\Gamma \vdash_{\text{smt}} \text{erase}(\tau)$.*

PROOF. By induction on the well formedness derivation.

(t -BASE) Immediate, since $\text{erase}(B) = B$.

(t -TVAR) Contradictory—type variables aren't well typed at smt.

(t -ADT) By the IH on each constituent of $D \vec{\tau}_i$, and then by t -ADT.

(τ -SMT) Since $\text{erase}(t \text{ smt}) = \text{erase}(t)$, by the IH on $\Gamma \vdash_{\text{smt}} t$.

(τ -SYM) Since $\text{erase}(t \text{ sym}) = \text{erase}(t)$, by the IH on $\Gamma \vdash_{\text{smt}} t$.

(τ -MODEL) Contradictory—model isn't well typed at smt. □

LEMMA D.3 (SMT TYPES HAVE ONLY SMT PARTS). *If $\Gamma \vdash_{\text{smt}} \tau$, then all of τ 's subparts are also well formed at smt.*

PROOF. By induction on τ .

$(\tau = B)$ Immediate.

$(\tau = \alpha)$ Contradictory—type variables are only well formed at exp.

$(\tau = D \vec{\tau}_i)$ By the IH on each τ_i .

$(\tau = t \text{ smt})$ By the IH on t .

$(\tau = t \text{ sym})$ By the IH on t .

$(\tau = \text{model})$ Contradictory—model is only well formed at exp. \square

LEMMA D.4 (SMT CONVERSION IS WELL FORMED). *If $\Gamma \vdash_{\text{smt}} \tau$ then $\text{toSMT}(\tau) = t \text{ sym}$ or $t \text{ smt}$ such that $\Gamma \vdash_{\text{smt}} t$ (and so $\Gamma \vdash_{\text{smt}} \text{toSMT}(\tau)$).*

PROOF. By induction on the well formedness derivation.

$(t\text{-BASE})$ $\text{toSMT}(B) = B \text{ smt}$, which is well formed by $\tau\text{-SMT}$ and $t\text{-B}$.

$(t\text{-TVAR})$ Contradictory—type variables are only well formed at exp.

$(t\text{-ADT})$ We know that $\text{erase}(D \vec{\tau}_i)$ is still well formed by Lemma D.2; then by $\tau\text{-SMT}$.

$(\tau\text{-SMT})$ Since it must be that $\Gamma \vdash_{\text{smt}} t$, then $\text{erase}(t)$ is also well formed by Lemma D.2; then by $\tau\text{-SMT}$.

$(\tau\text{-SYM})$ Since it must be that $\Gamma \vdash_{\text{smt}} t$, then $\text{erase}(t)$ is also well formed by Lemma D.2; then by $\tau\text{-SYM}$.

$(\tau\text{-MODEL})$ Contradictory—model is only well formed at smt \square

LEMMA D.5 (SMT CONVERSION IS ONLY FOR SMT TYPES). *$\text{toSMT}(\tau)$ is defined iff $\Gamma \vdash_{\text{smt}} \tau$.*

PROOF. The right-to-left direction is proved by Lemma D.4. For left-to-right, we go by induction on τ .

$(\tau = B)$ By $t\text{-BASE}$.

$(\tau = \alpha)$ Contradictory—type variables are undefined for erase.

$(\tau = D \vec{\tau}_i)$ By the IH on each τ_i and $t\text{-ADT}$.

$(\tau = t \text{ smt})$ By the IH on t and $\tau\text{-SMT}$.

$(\tau = t \text{ sym})$ By the IH on t and $\tau\text{-SYM}$.

$(\tau = \text{model})$ Contradictory—model is undefined for erase. \square

We say a type t is an ‘‘SMT type’’ when $\Gamma \vdash_{\text{smt}} t$; a type τ is an SMT type when it is equal to an SMT type t or when it is of the form $t \text{ smt}$ or $t \text{ sym}$. Note that toSMT always produces an SMT type, but does not work on types that contain type variables or the unrepresentable model type.

LEMMA D.6 (WEAKENING). *If $\vdash \Gamma$ and $\vdash \Gamma'$ and $\text{dom}(\Gamma) \cap \text{dom}(\Gamma') = \emptyset$ then:*

(1) $\vdash \Gamma, \Gamma'$

(2) *If $\Gamma \vdash_m \tau$ then $\Gamma, \Gamma' \vdash_m \tau$;*

(3) *If $\Gamma \vdash e : \tau$ then $\Gamma, \Gamma' \vdash e : \tau$; and*

(4) *If $\Gamma \vdash \phi : \tau$ then $\Gamma, \Gamma' \vdash \phi : \tau$.*

PROOF. By mutual induction on the derivations.

Contexts.

$(\Gamma\text{-EMPTY})$ We have $\Gamma' = \cdot$; immediate by assumption.

$(\Gamma\text{-VAR})$ We have $\Gamma' = \Gamma'', X : \tau$. By the IH on Γ'' and $\Gamma\text{-VAR}$, finding $\Gamma, \Gamma'' \vdash_m \tau$ by part (2) of the IH.

$(\Gamma\text{-TVAR})$ We have $\Gamma' = \Gamma'', \alpha$. By the IH on Γ'' and $\Gamma\text{-TVAR}$.

Type well formedness.

- (*t*-BASE) Immediate, by *t*-BASE.
- (*t*-TVAR) Since Γ and Γ' have disjoint domains, we know $\alpha \in \Gamma$ —by *t*-TVAR.
- (*t*-ADT) By the IH on each constituent of $D \vec{\tau}_i$, followed by *t*-ADT.
- (τ -SMT) By the IH on $\Gamma \vdash_{\text{smt}} t$ and then τ -SMT.
- (τ -SYM) By the IH on $\Gamma \vdash_{\text{smt}} t$ and then τ -SYM.
- (τ -MODEL) Immediate, by τ -MODEL. □

Expressions.

- (*e*-VAR) Since the domains are disjoint, $(\Gamma, \Gamma')(X) = \tau$ and we can still find *e*-VAR.
- (*e*-CONST) Immediate, by *e*-CONST.
- (*e*-LET) By the *e*-LET and the IH on e_1 and e_2 , α -renaming X appropriately.
- (*e*-CTOR) By *e*-CTOR and the IH, using part (2) on τ'_j and part (3) on e_i .
- (*e*-QUOTE) By the part (4) of the IH.
- (*e*-REL) By *e*-REL and the IH on each e_i .
- (*e*-FUN) By *e*-FUN and the the IH, using part (2) on τ'_j and part (3) on e_i .
- (*e*-OP) By *e*-OP and the the IH, using part (2) on τ'_j and part (3) on e_i .
- (*e*-IF) By *e*-IF and the IH on each of the e_i .
- (*e*-MATCH) By *e*-MATCH and the IH on e and each of the e_i , α -renaming each X_k appropriately.

Formulas.

- (ϕ -VAR) By ϕ -VAR and part (2) of the IH.
- (ϕ -PROMOTE) By ϕ -PROMOTE and the IH.
- (ϕ -UNQUOTE) By ϕ -UNQUOTE and part (3) of the IH.
- (ϕ -CTOR) By ϕ -CTOR and the IH, using part (2) on the τ_i and part (4) on ϕ_i , observing that the actual types are unchanged, and so the toSMT conversions are the same.

LEMMA D.7 (TYPE WELL FORMEDNESS STRENGTHENING). *If $\Gamma, X : \tau, \Gamma' \vdash_m \tau'$ then $\Gamma, \Gamma' \vdash \tau'$.*

PROOF. (*t*-BASE) Immediate.

- (*t*-TVAR) Immediate: removing the variable binding can't affect α .
- (*t*-ADT) By the IH on each constituent of $D \vec{\tau}_i$.
- (τ -SMT) By the IH on $\Gamma \vdash_{\text{smt}} t$.
- (τ -SYM) By the IH on $\Gamma \vdash_{\text{smt}} t$.
- (τ -MODEL) Immediate. □

LEMMA D.8 (TYPE VARIABLE SUBSTITUTION). *If $\vdash \Gamma, \alpha, \Gamma'$ and $\Gamma \vdash_m \tau'$, then:*

- (1) $\vdash \Gamma, \Gamma'[\tau/\alpha]$;
- (2) *If $\Gamma, \alpha, \Gamma' \vdash_m \tau$ then $\Gamma, \Gamma'[\tau/\alpha] \vdash_m \tau'[\tau/\alpha]$;*
- (3) *If $\Gamma, \alpha, \Gamma' \vdash X, \tau' \triangleright \Gamma''$ then $\Gamma, \Gamma'[\tau/\alpha] \vdash X, \tau'[\tau/\alpha] \triangleright \Gamma''[\tau/\alpha]$; and*
- (4) *If $\Gamma, \alpha, \Gamma' \vdash \vec{X}_i, \vec{\tau}'_i \triangleright \Gamma''$ then $\Gamma, \Gamma'[\tau/\alpha] \vdash \vec{X}_i, \vec{\tau}'_i[\tau/\alpha] \triangleright \Gamma''[\tau/\alpha]$.*

PROOF. For parts (1) and (2), by mutual induction on the derivations. Note that if α actually occurs in the type, we could only have found well formedness at exp.

(Γ -EMPTY) Contradictory: $\cdot \neq \Gamma, \alpha, \Gamma'$.

(Γ -VAR) We have $\Gamma' = \Gamma'', X : \tau'$ where $\Gamma, \Gamma'' \vdash_{\text{exp}} \tau'$. By the IH on Γ'' , we know that $\vdash \Gamma, \Gamma''[\tau/\alpha]$; by part (2), we have $\Gamma, \Gamma''[\tau/\alpha] \vdash_{\text{exp}} \tau'[\tau/\alpha]$; and so we have $\vdash \Gamma, \Gamma'[\tau/\alpha]$ by Γ -VAR.

(Γ -TVAR) We have $\Gamma' = \Gamma'', \beta$; by the IH on Γ'' , we have $\vdash \Gamma, \Gamma''[\tau/\alpha]$; since $\beta[\tau/\alpha] = \beta$, we can apply Γ -TVAR to find $\vdash \Gamma, \Gamma'[\tau/\alpha]$ as desired.

(t -B) Immediate by t -B, since $B[\tau/\alpha] = B$.

(t -TVAR) We have $\tau' = \beta$. If $\alpha = \beta$, then we have $\Gamma \vdash_m \alpha[\tau/\alpha]$ by assumption. If $\alpha \neq \beta$, then it must be that $\beta \in \Gamma$ or Γ' —either way, β is unaffected by the substitution and we have $\beta \in \Gamma, \Gamma'[\tau/\alpha]$ and so $\Gamma, \Gamma'[\tau/\alpha] \vdash_m \beta$ by t -TVAR.

(t -ADT) By the IH on each premise, followed by t -ADT.

(τ -SMT) By the IH on $\Gamma, \alpha, \Gamma' \vdash_{\text{smt}} t$ and then by τ -SMT.

(τ -SYM) By the IH on $\Gamma, \alpha, \Gamma' \vdash_{\text{smt}} t$ and then by τ -SYM.

(τ -MODEL) Immediate by τ -MODEL.

For parts (3) and (4), by mutual induction on the derivations.

($X\tau$ -BIND) We have $X \notin \text{dom}(\Gamma, \alpha, \Gamma')$, so it must also be the case that $X \notin \text{dom}(\Gamma, \Gamma'[\tau/\alpha])$. We therefore find $\Gamma, \Gamma'[\tau/\alpha] \vdash X, \tau'[\tau/\alpha] \triangleright \Gamma, \Gamma'[\tau/\alpha], \tau'[\tau/\alpha]$ by $X\tau$ -BIND.

($X\tau$ -CHECK) We have $(\Gamma, \alpha, \Gamma')(X) = \tau'$. Is $X : \tau'$ in Γ or Γ' ? Either way we will find $\Gamma, \Gamma'[\tau/\alpha] \vdash X, \tau'[\tau/\alpha] \triangleright \Gamma, \Gamma'[\tau/\alpha]$ by $X\tau$ -CHECK.

If $\tau' \in \text{dom}(\Gamma)$, then $\Gamma \vdash \tau'$ and so $\tau'[\tau/\alpha] = \tau'$ ($\Gamma, \Gamma'[\tau/\alpha])(X) = \tau'$ and we have $\Gamma, \Gamma'[\tau/\alpha] \vdash X, \tau' \triangleright \Gamma, \Gamma'[\tau/\alpha]$.

If, on the other hand, $\tau' \in \text{dom}(\Gamma')$, then $(\Gamma, \Gamma'[\tau/\alpha])(X) = \tau'[\tau/\alpha]$. We therefore have $\Gamma, \Gamma'[\tau/\alpha] \vdash X, \tau'[\tau/\alpha] \triangleright \Gamma, \Gamma'[\tau/\alpha]$.

($\vec{X}\vec{\tau}$ -ALL) By part (3) of the IH on each premise. □

LEMMA D.9 (REGULARITY; FORMULAS HAVE SMT TYPES). (1) If $\vdash \Gamma$ and $\Gamma(X) = \tau$ then $\Gamma \vdash_{\text{exp}} \tau$.

(2) If $\vdash \Phi$ then (a) if $f : \forall \vec{\alpha}_j, \vec{\tau}_i \rightarrow \tau \in \Phi$ then $\vec{\alpha}_j \vdash_{\text{exp}} \tau$, and (b) if $uf : \vec{t}'_i \rightarrow t \in \Phi$ then $\cdot \vdash_{\text{smt}} t$.

(3) If $\Delta; \Phi; \Gamma \vdash e : \tau$ then $\Gamma \vdash_{\text{exp}} \tau$.

(4) If $\Delta; \Phi; \Gamma \vdash c_c^{\text{SMT}}(\vec{\phi}_i) : \vec{\tau}_i \rightarrow \tau$ then $\tau_i = t_i$ smt or $\tau_i = t_i$ sym and $\tau = t$ smt or $\tau = t$ sym and $\Gamma \vdash_{\text{smt}} \tau_i$ and $\Gamma \vdash_{\text{smt}} \tau$.

(5) If $\Delta; \Phi; \Gamma \vdash \phi : \tau$ then $\tau = t$ smt or $\tau = t$ sym and $\Gamma \vdash_{\text{smt}} \tau$.

(6) If $\Delta; \Phi; \Gamma \vdash \vec{e}_i : \vec{\tau}_i$ then $\Gamma \vdash_{\text{exp}} \tau_i$.

(7) If $\Delta; \Phi; \Gamma \vdash \vec{\phi}_i : \vec{\tau}_i$ then $\tau_i = t_i$ smt or $\tau = t_i$ sym and $\Gamma \vdash_{\text{smt}} t_i$ (and so $\Gamma \vdash_{\text{exp}} \tau_i$).

PROOF. By induction on the typing derivation.

Contexts.

(Γ -EMPTY) Contradictory—there's no way \cdot has a binding for X .

(Γ -VAR) $\Gamma = \Gamma', Y : \tau$. If $X = Y$, then we know $\Gamma \vdash_{\text{exp}} \tau$ by assumption; otherwise, by the IH on Γ' .

(Γ -TVAR) $\Gamma = \Gamma', \alpha$. By the IH on Γ .

Program signatures.

(Φ -EMPTY) Contradictory—there are no function definitions in \cdot .

(Φ -FUN) $\Phi = \Phi', g : \dots$ For case (a) when $f = g$, then by assumption. Otherwise, by the IH on Φ' .

(Φ -REL) $\Phi = \Phi', p \subseteq \vec{\tau}_i$. By the IH on Φ .

(Φ -UFUN) $\Phi = \Phi', uf' : \vec{t}_i \rightarrow t$. For case (b) when $uf = uf'$, then by assumption. Otherwise, by the IH on Φ' .

Expressions.

(e -VAR) By the part (1) on $\vdash \Gamma$.

(e -CONST) By assumption, we know that $\Gamma \vdash_{\text{smt}} \text{typeof}(k)$; by Lemma D.1 we can find $\Gamma \vdash_{\text{exp}} \text{typeof}(k)$.

(e -LET) By the IH on $\Gamma, X : \tau_1 \vdash e_2 : \tau_2$, using strengthening (Lemma D.7) to find that if $\Gamma, X : \tau_1 \vdash_{\text{exp}} \tau_2$ then $\Gamma \vdash_{\text{exp}} \tau_2$.

(e -CTOR) Since $\Gamma \vdash_{\text{exp}} \tau'_j$, we know by t -ADT that $\Gamma \vdash_{\text{exp}} D \vec{\tau}'_j$.

(e -QUOTE) By the IH on part (5), we know that $\Gamma \vdash_{\text{smt}} \tau$ (and, less relevantly, that $\tau = t$ smt or t sym). We can find the same well formedness at exp by Lemma D.1.

(e -REL) Immediate by t -B.

(e -FUN) Since $f : \forall \vec{\alpha}_j, \vec{\tau}_i \rightarrow \tau \in \Phi$ and $\vdash \Phi$, we know by part (2) of the IH know that $\vec{\alpha}_j \vdash_{\text{exp}} \tau$. By weakening (Lemma D.6) we can lift that well formedness judgment to Γ . Since each $\Gamma \vdash_{\text{exp}} \tau'_j$, we can find that $\Gamma \vdash \tau[\tau'_j/\alpha_j]$ by substitution (Lemma D.8).

(e -FUN) We have by assumption that $\text{typeof}(\otimes)$ yields a well-formed type, i.e., $\vec{\alpha}_j \vdash_{\text{exp}} \tau$ (and also for each τ_i). By weakening (Lemma D.6) we can lift that well formedness judgment to Γ . Since each $\Gamma \vdash_{\text{exp}} \tau'_j$, we can find that $\Gamma \vdash \tau[\tau'_j/\alpha_j]$ by substitution (Lemma D.8).

(e -IF) By the IH on $\Gamma \vdash e_2 : t$.

(e -MATCH) By the IH on $\Gamma, \overline{X_1 : \tau_1[\tau_j/\alpha_j]} \vdash e_1 : \tau$ we have $\Gamma, \overline{X_1 : \tau_1[\tau_j/\alpha_j]} \vdash \tau$; we can use strengthening (Lemma D.7) to find $\Gamma \vdash_{\text{exp}} \tau$.

SMT constructors.

(c -SMT-VAR) Immediate, with $\Gamma \vdash_{\text{smt}} t$ coming from the rule itself.

(c -SMT-CONST) Immediate, since we have by assumption that $\cdot \vdash_{\text{smt}} \text{typeof}(k)$.

(c -SMT-LET) Immediate, with the necessary well formedness assumptions coming from the rule itself.

(c -SMT-CTOR) We know that $D \vec{t}'_j$ is well formed by t -ADT. We can find the translation of the argument types well formed by Lemma D.4 on each of the $\Gamma \vdash_{\text{smt}} \tau_i[t'_j/\alpha_j]$ derivations.

(c -SMT-FORALL) Immediate, with the necessary $\Gamma \vdash_{\text{smt}} t_1$ coming from the rule itself.

(c -SMT-UFUN) Since $\vdash \Phi$ and $uf : \vec{t}_i \rightarrow t \in \Phi$, we know that $\cdot \vdash_{\text{smt}} t$ by part (2) of the IH and so $\cdot \vdash_{\text{smt}} t$ smt, which we can lift to Γ by weakening (Lemma D.6).

Formulas.

(ϕ -PROMOTE) Since $\Delta; \Phi; \Gamma \vdash \phi : t$ sym, we know that $\Gamma \vdash_{\text{smt}} t$ and so we are correct in yielding $\tau = t$ smt.

(ϕ -UNQUOTE) We have $\Delta; \Phi; \Gamma \vdash e : \tau$ such that $\Gamma \vdash_{\text{smt}} \tau$. By Lemma D.4 we know that $\text{toSMT}(\tau)$ is a well formed SMT type.

(ϕ -CTOR) By the IH part (4) on $\Gamma \vdash c_c^{\text{SMT}} : \vec{\tau}_i \rightarrow \tau$.

Vectored expressions and formulas. By the IH for parts (3) and (5), respectively. \square

LEMMA D.10 (SMT VALUE CONVERSION IS TYPE CORRECT). *If $\Delta; \Phi; \Gamma \vdash v : \tau$ and $\Gamma \vdash_{\text{smt}} \tau$ then $\Delta; \Phi; \Gamma \vdash \text{toSMT}(v) : \text{toSMT}(\tau)$.*

PROOF. First, observe that τ and all of its parts must be well formed at smt, by Lemmas D.5 and D.3. By induction on the typing derivation. In expression mode, the applicable rules are e -CONST and e -CTOR; only a few typing rules could even have applied to a value in formula mode: the ϕ -CTOR and ϕ -PROMOTE.

(e -CONST) We have $\Gamma \vdash k : \text{typeof}(k)$; since $\text{toSMT}(k) = c_{\text{const}}^{\text{SMT}}[k]()$ and $\text{toSMT}(\text{typeof}(k)) = k$ smt (since $\Gamma \vdash_{\text{smt}} \text{typeof}(k)$ by assumption), we must show that $\Gamma \vdash c_{\text{const}}^{\text{SMT}}[k]() : \text{typeof}(k)$ smt, which we have by ϕ -SMT-CONST.

(e -CTOR) We have $v = c(\vec{v}_i)$ and:

$$\Delta(D) = \forall \vec{\alpha}_j, \{ \dots, c : \vec{\tau}_i, \dots \} \quad \Gamma \vdash_{\text{exp}} \tau'_j \quad \Gamma \vdash v_i : \tau_i[\tau'_j / \alpha_j]$$

Further, we know that $\Gamma \vdash_{\text{smt}} D \vec{\tau}'_j$ (and so each of the subderivations must also be smt) and that $\text{toSMT}(c(\vec{v}_i)) = c_{\text{ctor}}^{\text{SMT}}[c](\overrightarrow{\text{toSMT}(v_i)})$. By the IH on each of these v_i , we know that we find appropriate values at appropriately converted types, i.e., $\Gamma \vdash \text{toSMT}(v_i) : \text{toSMT}(\tau_i[\tau'_j / \alpha_j])$. By Lemma D.4, we know $\text{toSMT}(\tau_i[\tau'_j / \alpha_j])$ is some well formed SMT type. We are *almost* able to apply ϕ -SMT-CTOR, but we must pick appropriate t'_j . We know that $\text{toSMT}(\tau'_j)$ is a well formed SMT type of the form t'_j smt or t'_j sym (Lemma D.4). Whether it's symbolic or not, let the inner t'_j there be our t'_j . We can now apply ϕ -SMT-CTOR to find that $\Gamma \vdash \text{toSMT}(c(\vec{v}_i)) : \text{toSMT}(D \vec{\tau}'_j)$.

(ϕ -PROMOTE) By the IH on $\Gamma \vdash v : t$ sym, we know that $\Gamma \vdash \text{toSMT}(v) : \text{toSMT}(t$ sym), i.e., $\Gamma \vdash \text{toSMT}(v) : t$ sym. By reapplying ϕ -PROMOTE we can find that $\Gamma \vdash \text{toSMT}(v) : v$ smt.

(ϕ -CTOR) Immediate: toSMT does nothing to the c_{\dots}^{SMT} constructed value nor to the SMT-type τ assigned to it (which is t smt in all cases except for $c_{\text{var}}^{\text{SMT}}$). \square

E TYPE SAFETY

To prove type safety, we prove two properties for every mode of evaluation: first, it is *safe*, i.e. never yields \perp ; and second, it is *type preserving*, i.e., well typed inputs yield well typed outputs.

The proofs are fairly conventional. For all but the last step, we prove safety and type preservation simultaneously. We start with expressions and formulas (Lemma E.2), which requires a modest notion of canonical forms (Lemma E.1). Next, we prove that value unification (Lemma E.3) is type preserving, reasoning about unification in general within the lemma showing safety and type preservation for premises (Lemma E.4). After a brief lemma about bindings (Lemma E.5), we can prove that program evaluation is safe (Theorem E.6) and type preserving (Theorem E.7).

LEMMA E.1 (CANONICAL FORMS FOR t sym). *If $\Delta; \Phi; \Gamma \vdash v : t$ sym then $v = c_{\text{var}}^{\text{SMT}}[x, t]()$.*

PROOF. The only typing rule that could have applied is ϕ -SMT-VAR. \square

LEMMA E.2 (TERM AND FORMULA TYPE SAFETY). *If $\Delta; \Phi \models \mathcal{W}$ and $\Delta; \Phi \vdash \vec{F}$ and $\Gamma \models \theta$, when either:*

- (1) $\Delta; \Phi; \Gamma \vdash e : \tau$ and $\mathcal{W}; \theta \vdash e \Downarrow_e v_\perp$; or
- (2) $\Delta; \Phi; \Gamma \vdash \phi : \tau$ and $\mathcal{W}; \theta \vdash \phi \Downarrow_\phi v_\perp$
- (3) $\Delta; \Phi \vdash \text{fun } f(\vec{X}_i : \vec{\tau}_i) : \tau = e$ and $\vec{\alpha}_j, \vec{X}_i : \vec{\tau}_i \models \theta'$ and $\mathcal{W}; \theta' \vdash e \Downarrow_e v_\perp$

then $v_\perp = v$ (i.e., $v_\perp \neq \perp$) and $\Delta; \Phi; \Gamma \vdash v : \tau$.

Similarly, when either:

- (1) if $\Delta; \Phi; \Gamma \vdash e_i : \tau_i$ and $\mathcal{W}; \theta \vdash \vec{e}_i \Downarrow_e \vec{v}_i$ then $\Delta; \Phi; \Gamma \vdash \vec{v}_{i_\perp} : \vec{\tau}_i$; and
- (2) if $\Delta; \Phi; \Gamma \vdash \phi_i : \tau_i$ and $\mathcal{W}; \theta \vdash \vec{\phi}_i \Downarrow_\phi \vec{v}_i$ then $\Delta; \Phi; \Gamma \vdash \vec{v}_{i_\perp} : \tau$

then $\vec{v}_{i_\perp} = \vec{v}_i$ (i.e., it is not \perp) and $\Delta; \Phi; \Gamma \vdash v_i : \tau_i$.

PROOF. By mutual induction on derivations and the length of the vectored expressions/formulas, leaving θ general (for, e.g., e -LET and e -MATCH).

Expressions.

- (e -VAR) We have $\Gamma(X) = \tau$; since $\Gamma \models \theta$, we have $\theta(X) = v$ (and so \Downarrow_e -VAR-E didn't apply). So it must be the case that \Downarrow_e -VAR applied. We can see further that $\Delta; \Phi; \cdot \vdash v : \tau$, and we are done by weakening (Lemma D.6).
- (e -CONST) It must be that \Downarrow_e -CONST applied, and we immediately see that $v_\perp \neq \perp$ and k is well typed in any well formed context by assumption and e -CONST.
- (e -LET) We know that $\Gamma \vdash e_1 : \tau_1$ and $\Gamma, X : \tau_1 \vdash e_2 : \tau_2$. By the IH on e_1 , we know that $\theta; \mathcal{W} \vdash e_1 \Downarrow_e v_1$, so it can't be the case that \Downarrow_e -LET-E applied—it must have been \Downarrow_e -LET. By the IH on e_2 , we know that the final result is also not \perp and is well typed.
- (e -CTOR) We have $\Delta(D) = \forall \vec{\alpha}_j, \{ \dots, c : \vec{\tau}_i, \dots \}$ and $\Gamma \vdash e_i : \tau_i[\tau'_j/\alpha_j]$. By the IH, we know that each of the \vec{e}_i must have reduced to non- \perp values, and so \Downarrow_e -CTOR-E could not have applied. We can therefore see that each e_i reduces to an appropriately typed v_i , and our resulting value is well typed by e -CTOR.
- (e -QUOTE) Only \Downarrow_e -QUOTE could have applied. By the IH, we know that ϕ reduces to a non- \perp value v well typed at τ .
- (e -REL) We know that $p \subseteq \vec{\tau}_i \in \Phi$ and $\Gamma \vdash e_i : \tau_i$. The IH on \vec{e}_i rules out $StepstoE$ -REL-E1; the typing rule rules out the arity mismatch in \Downarrow_e -REL-E2 and the missing relation in \Downarrow_e -REL-E3. So it must be the case that \Downarrow_e -REL-TRUE or \Downarrow_e -REL-FALSE applied; either way, we yield a bool, which is appropriately typed by e -CONST.
- (e -FUN) We know that $f : \forall \vec{\alpha}_j, \vec{\tau}_i \rightarrow \tau \in \Phi$ and $\Gamma \vdash e_i : \tau_i[\tau'_j/\alpha_j]$. The IH on \vec{e}_i rules out $StepstoE$ -FUN-E1; the typing rule rules out the arity mismatch in \Downarrow_e -FUN-E2 and the missing function in \Downarrow_e -FUN-E3. So it must be the case that \Downarrow_e -FUN applied. Since $\Delta; \Phi \vdash F$, we know by the IH on part (3) that the resulting value is non- \perp and well typed at $\tau[\tau'_j/\alpha_j]$.
- (e -OP) We know that $\text{typeof}(\otimes) = \forall \vec{\alpha}_j, \vec{\tau}_i \rightarrow \tau$ and $\Gamma \vdash e_i : \tau_i[\tau'_j/\alpha_j]$. The IH on \vec{e}_i rules out $StepstoE$ -OP-E1; the typing rule rules out the arity/domain mismatch in \Downarrow_e -OP-E2. So it must be the case that \Downarrow_e -OP applied. We know that the result is well typed by our assumption that $\text{typeof}(\otimes)$ and $\llbracket \otimes \rrbracket$ agree.
- (e -IF) We have $\Gamma \vdash e_1 : \text{bool}$ and $\Gamma \vdash e_2 : \tau$ and $\Gamma \vdash e_3 : \tau$. By the IH on e_1 , we know that e_1 reduces to true or false (since those are the only values of type bool). So we

can rule out $\Downarrow_e\text{-ITE-E1}$ and $\Downarrow_e\text{-ITE-E2}$ —we must have stepped by either $\Downarrow_e\text{-ITE-T}$ or $\Downarrow_e\text{-ITE-F}$. The IH on e_2 or e_3 (respectively) guarantees we step to a non- \perp , well typed value.

(*e*-MATCH) We have $\Gamma \vdash e : D \vec{\tau}_j$ and $\Delta(D) = \forall \vec{\alpha}_j, \{ \dots, c_i : \vec{\tau}_k, \dots \}$ and $\Gamma, \overrightarrow{X_k : \tau_k[\tau_j/\alpha_j]} \vdash e_i : \tau$. The IH on e guarantees that we get a non- \perp value at type $D \vec{\tau}_j$, which rules out the error case $\Downarrow_e\text{-MATCH-E1}$, the non-constructor value of $\Downarrow_e\text{-MATCH-E2}$, the misnamed constructor of $\Downarrow_e\text{-MATCH-E3}$, and the arity error of $\Downarrow_e\text{-MATCH-E4}$. So it must be the case that we applied $\Downarrow_e\text{-MATCH}$; by the IH, the matching pattern reduces to a well typed non- \perp value.

Formulas.

(ϕ -PROMOTE) We have $\Delta; \Phi; \Gamma \vdash \phi : t \text{ sym}$; by the IH, we know that ϕ steps to a non- \perp value v well typed at $t \text{ sym}$; by ϕ -PROMOTE we can see that v is also well typed at $t \text{ smt}$.

(ϕ -UNQUOTE) We have $,e;$ since $\Delta; \Phi; \Gamma \vdash e : \tau$, we know by the IH that e reduces to a non- \perp value v that is also well typed at τ . We can therefore rule out $\Downarrow_\phi\text{-UNQUOTE-E}$, so we must have stepped by $\Downarrow_\phi\text{-UNQUOTE}$.

Since $\Delta; \Phi; \Gamma \vdash v : \tau$ and $\Gamma \vdash_{\text{smt}} \tau$, we have $\Delta; \Phi; \Gamma \vdash \text{toSMT}(v) : \text{toSMT}(\tau)$ by Lemma D.10, as desired.

(ϕ -CTOR) We have $c_c^{\text{SMT}}(\vec{\phi}_i)$ such that $\Gamma \vdash c_c^{\text{SMT}} : \vec{\tau}_i \rightarrow \tau$ and $\Gamma \vdash \phi_i : \tau_i$. We know by the IH that each ϕ_i is well typed at τ_i and so none of them step to \perp , and so $\Downarrow_\phi\text{-CTOR-E}$ cannot apply.

Therefore either $\Downarrow_\phi\text{-CTOR}$ or $\Downarrow_\phi\text{-VALUE}$ applied; the resulting value is well typed by the IH or remains well typed, respectively.

Functions. By part (1) on $\vec{\alpha}_j, \overrightarrow{X_i : \tau_i} \vdash e : \tau$, using weakening (Lemma D.6) to recover typing in Γ .

Vectored expressions and formulas. By induction on the vector length, using parts (1) and (2) in each case. \square

LEMMA E.3 (VALUE UNIFICATION PRESERVATION). *If $\Gamma \models \theta$ when either:*

- (1) $\Gamma \vdash \vec{X}, \vec{\tau} \triangleright \Gamma'$ and $\Gamma \vdash \vec{v} : \vec{\tau}$ and $\theta \vdash \vec{X} \sim \vec{v} : \theta'$; or
- (2) $\Gamma \vdash X, \tau \triangleright \Gamma'$ and $\Gamma \vdash v : \tau$ and $\theta \vdash X \sim v \triangleright \theta'$;

then $\Gamma' \models \theta'$.

PROOF. By induction on the derivation of well typing.

($X\tau$ -BIND) Only $uv\text{-BIND-VAR}$ could have applied, so we have $\Gamma \vdash v : \tau$ and $\Gamma \models \theta$ and must show that $\Gamma, X : \tau \models \theta[X \mapsto v]$, which we have immediately.

($X\tau$ -CHECK) Here $X \in \Gamma$, so it must be that $\theta(X)$ is defined. One of three rules could have applied:

(uv -EQ-VAR) We have $\Gamma' = \Gamma$ and $\theta' = \theta$, so $\Gamma' \models \theta'$ by assumption.

(uv -CTOR) By the IH, we know that Γ' models θ' .

(uv -CONSTANT) As for $uv\text{-EQ-VAR}$, we have $\Gamma' = \Gamma$ and $\theta' = \theta$, so $\Gamma' \models \theta'$ by assumption.

($X\tau$ -ALL) It must be that $\vec{u}\vec{v}\text{-ALL}$ applied; by the IH on each sub-derivation, we can find that $\Gamma_i \models \theta_i$, and so $\Gamma' \models \theta'$ in particular. \square

User code will never *directly* trigger a use of uv -EQ-VAR directly, because the unification rules won't call value unification with a defined LHS (we'd just use uu -BB instead). But a use of $\vec{u}\vec{v}$ -ALL could lead to a variable being unified early on and then used again in the same unification process.

LEMMA E.4 (PREMISE PRESERVATION AND SAFETY). *If $\Delta; \Phi; \Gamma \vdash P \triangleright \Gamma'$ and $\Delta; \Phi \models \vec{F}$ and $\Delta; \Phi \models \mathcal{W}$ and $\Gamma \models \theta$ then if $\vec{F}; \mathcal{W}; \theta \vdash P \rightarrow \theta'_\perp$ then:*

- (1) $\theta'_\perp = \theta'$ (i.e., it is not \perp); and
- (2) $\Gamma' \models \theta'$.

PROOF. By induction on the premise typing derivation, followed by cases on the step taken.

(P-PosATOM) We have:

$$p \subseteq \vec{\tau}_i \in \Phi \quad \Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'$$

The only rule that could have applied is PosATOM, i.e., $\vec{v} \in \mathcal{W}(p)$ and $\theta \vdash \vec{X}_i \sim \vec{v}_i : \theta'_\perp$. We must show that $\theta'_\perp = \theta'$ and $\Gamma' \models \theta'$.

Since $\Delta; \Phi \models \mathcal{W}$, we know that $\cdot \vdash \vec{v}_i : \vec{\tau}_i$; by weakening we have $\Gamma \vdash \vec{v}_i : \vec{\tau}_i$ (Lemma D.6).

Syntactically, we know that \vec{X}_i are all variables and that \vec{v}_i are all values. For each one, therefore only two unification rules could possibly apply: uu -BB (X_i is bound) and uu -FB (X_i is free). In particular, uu -FF cannot apply, and so we cannot produce \perp , so $\theta'_\perp = \theta' = \theta\vec{\theta}_i$. By Lemma E.3, we know that $\Gamma' \models \theta\vec{\theta}_i$ for each i , and so $\Gamma' \models \theta\vec{\theta}_i$.

(P-NEGATOM) We have:

$$p \subseteq \vec{\tau}_i \in \Phi \quad \Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma$$

Two rules are possible: NEGATOM and NEGATOM-E. We must show that the latter cannot apply and that the former preserves typing.

Since $\Gamma \models \theta$, it must be that case that each $\vec{X}_i \in \text{dom}(\theta)$, and so NEGATOM-E cannot have applied. It remains to be seen that $\Gamma \models \theta'$ —but in NEGATOM we have $\theta = \theta'$, and so we are done.

(P-EQCTOR-BF) We have:

$$\Delta(D) = \forall \vec{\alpha}_j, \{ \dots, c : \vec{\tau}_i, \dots \} \\ \Gamma \vdash Y, \tau[\vec{\tau}'_j / \vec{\alpha}_j] \triangleright \Gamma \quad \vec{X}_i \not\subseteq \Gamma \quad \Gamma \vdash \vec{X}_i, \vec{\tau}_i[\vec{\tau}'_j / \vec{\alpha}_j] \triangleright \Gamma'$$

The only rule that could have applied is EQCTOR, where $\theta \vdash Y c(\vec{X}_i) : \theta'_\perp$. We must show that $\theta'_\perp = \theta'$ (i.e., it is not \perp) and that $\Gamma' \models \theta'$.

Since $\Gamma \vdash Y, \tau[\vec{\tau}'_j / \vec{\alpha}_j] \triangleright \Gamma$, it must be the case that $Y \in \text{dom}(\Gamma)$ and so $\theta(Y) = v$ (and so $\cdot \vdash v : \tau[\vec{\tau}'_j / \vec{\alpha}_j]$, which also holds under Γ thanks to weakening (Lemma D.6)).

Only two rules could have applied to show $\theta \vdash Y c(\vec{X}_i) : \theta'_\perp$: uu -BF (when some of \vec{X}_i are unbound) or uu -BB (when all of the \vec{X}_i are bound). In either case, uu -FF can't have a applied, and so $\theta'_\perp = \theta'$.

One of two rules could have applied: uv -EQ-VAR or uv -CTOR.

In the former case, we applied uu -BB, because $\theta(c(\vec{X}_i)) = c(\vec{v}_i)$. We have $\theta' = \theta[X \mapsto c(\vec{v}_i)]$ and $\Gamma' \models \theta'$ by substitution on $\Gamma \vdash \vec{X}_i, \vec{\tau}_i[\vec{\tau}'_j / \vec{\alpha}_j] \triangleright \Gamma'$ (Lemma D.8).

In the latter case, we can find that $\Gamma' \models \theta'$ by Lemma E.3 on the assumption that $\Gamma \vdash \vec{X}_i, \vec{\tau}_i[\vec{\tau}'_j / \vec{\alpha}_j] \triangleright \Gamma'$, and the fact $\theta(Y) = v$ is well typed in Γ .

(*P*-EQSMT-BF) We have:

$$\Gamma \vdash c_{\mathcal{C}}^{\text{SMT}} : \vec{\tau}_i \rightarrow \tau \quad \Gamma \vdash Y, \tau \triangleright \Gamma \quad \vec{X}_i \not\subseteq \Gamma \quad \Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'$$

The only rule that could have applied is EQSMT, where $\theta \vdash Y \sim c_{\mathcal{C}}^{\text{SMT}}(\vec{X}_i) : \theta'_{\perp}$. We must show that $\theta'_{\perp} = \theta'$ (i.e., it is not \perp and that $\Gamma \models \theta'$).

Since $\Gamma \vdash Y, \tau \triangleright \Gamma$, it must be the case that $Y \in \text{dom}(\Gamma)$ and so $\theta(Y) = v$. We can conclude that $\cdot \vdash v : \tau$ and so $\Gamma \vdash v : \tau$ (Lemma D.6).

Only two rules could have applied to show $\theta \vdash Y c_{\mathcal{C}}^{\text{SMT}}(\vec{X}_i) : \theta'_{\perp}$, noting the removal of the unquote, since unification doesn't care: *uu*-BF (when some of \vec{X}_i are unbound) or *uu*-BB (when all of the \vec{X}_i are bound). In either case, *uu*-FF can't have applied, and so $\theta'_{\perp} = \theta'$.

It remains to show that $\Gamma \models \theta'$. If the outer unification rule was *uu*-BB, we have $\theta = \theta'$ and so $\Gamma \models \theta'$ by assumption. If outer unification rule was *uu*-BF, one of two rules could have applied to find value unification: either *uv*-EQ-VAR or *uv*-CTOR.

In the former case, we apply *uu*-BB inside, because $\theta(Y) = c_{\mathcal{C}}^{\text{SMT}}(\vec{v}_i)$. We have $\theta' = \theta[X \mapsto c_{\mathcal{C}}^{\text{SMT}}(\vec{v}_i)]$ and $\Gamma' \models \theta'$ by substitution on $\Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'$ (Lemma D.8).

In the latter case, we can find that $\Gamma' \models \theta'$ by Lemma E.3 on the assumption that $\Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma'$, and the fact $\theta(Y) = v$ is well typed in Γ .

(*P*-EQ-FB) We have:

$$\Gamma \vdash e : \tau \quad \Gamma \vdash Y, \tau \triangleright \Gamma'$$

The two possible rules are EQEXPR and EQEXPR-E. We must show that the latter could not have applied (and so $\theta'_{\perp} = \theta'$) and that $\Gamma' \models \theta'$. By Lemma E.2, we know that EQEXPR-E cannot apply and that $\mathcal{W}; \theta \vdash e \Downarrow_e v$ (and so $\Gamma \vdash v : \tau$).

Since v is a value, either *uu*-FB or *uu*-BB applied, depending on whether or not Y is bound. Either way, *uu*-FF couldn't have applied, and so $\theta'_{\perp} = \theta'$.

We can find that $\Gamma' \models \theta'$ by Lemma E.3 on $\Gamma \vdash Y, \tau \triangleright \Gamma'$ (along with the well typing of v).

(*P*-NEG EQ) We have $\Gamma \vdash e : \tau$ and $\Gamma \vdash Y, \tau \triangleright \Gamma$. By Lemma E.2, we know that NEG EQEXPR-E1 cannot apply and that $\mathcal{W}; \theta \vdash e \Downarrow_e v$ (and so $\Gamma \vdash v : \tau$). Since $\Gamma \models \theta$, it must be that case that $Y \in \text{dom}(\theta)$, and so NEGATOM-E2 cannot have applied. It remains to be seen that $\Gamma \models \theta'$ —but in NEGEXPR we have $\theta = \theta'$, and so we are done. \square

LEMMA E.5 (IDENTICAL BINDINGS IMPLIES CONTAINMENT). *If $\Gamma \vdash X, \tau \triangleright \Gamma$, then $X \in \text{dom}(\Gamma)$. Similarly, if $\Gamma \vdash \vec{X}_i, \vec{\tau}_i \triangleright \Gamma$, then $\vec{X}_i \subseteq \text{dom}(\Gamma)$.*

PROOF. By induction on the derivation.

(*X* τ -BIND) Contradictory: this rule could not have applied, since $\Gamma \neq \Gamma, X : \tau$.

(*X* τ -CHECK) We have $X \in \text{dom}(\Gamma)$ by assumption.

($\vec{X}\vec{\tau}$ -ALL) By the IH on each of our premises. \square

THEOREM E.6 (PROGRAM SAFETY). *If $\Delta; \Phi \vdash \vec{F}_i \vec{H}_j$ and $\Delta; \Phi \models \mathcal{W}$ then for all $H \in \vec{H}_j$, $\neg(\vec{F}_i; \mathcal{W} \vdash H \rightarrow \perp)$.*

PROOF. The program $\text{prog} = \vec{F}_i \vec{H}_j$ must have been well typed according to prog-WF, and so we have $\vdash \Delta$ and $\vdash \Phi$ along with derivations for each F and H :

$$\begin{array}{ccccccc} \Delta; \Phi \vdash F_0 & \dots & \Delta; \Phi \vdash F_i & \dots & \Delta; \Phi \vdash F_n \\ \Delta; \Phi \vdash H_0 & \dots & \Delta; \Phi \vdash H_j & \dots & \Delta; \Phi \vdash H_m \end{array}$$

Let an $H = p(X_k) :- \vec{P}_\ell \in \vec{H}_j$ be given. We know that $\Delta; \Phi \vdash H$ by H -CLAUSE, i.e.:

$$\begin{array}{ccccccc} \cdot \vdash P_0 \triangleright \Gamma_1 & \dots & \Gamma_\ell \vdash P_\ell \triangleright \Gamma_{\ell+1} & \dots & \Gamma_p \vdash P_p \triangleright \Gamma' \\ p \subseteq \vec{\tau}_k \in \Phi & & \Gamma' \vdash \vec{X}_k, \vec{\tau}_k \triangleright \Gamma' & & \end{array}$$

Let \mathcal{W} be given such that $\Delta; \Phi \models \mathcal{W}$. We must show that it is not the case that $\vec{F}_i; \mathcal{W} \vdash H \rightarrow \perp$, i.e., CLAUSE-E1 and CLAUSE-E2 cannot apply. We can rule out CLAUSE-E1 by Lemma E.4(1): it is not the case that a typesafe premise steps to \perp . To rule out CLAUSE-E2, we need to know that if we can build a final substitution, i.e.:

$$\cdot \vdash P_0 \rightarrow \theta_1 \quad \dots \quad \theta_\ell \vdash P_\ell \rightarrow \theta_{\ell+1} \quad \dots \quad \theta_p \vdash P_p \rightarrow \theta$$

then $\vec{X}_k \in \text{dom}(\theta)$. We know that $\vec{X}_k \subseteq \text{dom}(\Gamma')$ by Lemma E.5 on $\Gamma' \vdash \vec{X}_k, \vec{\tau}_k \triangleright \Gamma'$; since $\Gamma_p \vdash P_p \triangleright \Gamma'$, we know by Lemma E.4(2) that $\Gamma' \models \theta$. We can therefore conclude that $\forall X \in \text{dom}(\Gamma'), X \in \text{dom}(\theta)$, and so $\vec{X}_k \in \text{dom}(\theta)$... and CLAUSE-E2 cannot apply. \square

THEOREM E.7 (PROGRAM PRESERVATION). *If $\Delta; \Phi \vdash \vec{F}_i \vec{H}_j$ and $\Delta; \Phi \models \mathcal{W}$ and $\vec{F}_i; \mathcal{W} \vdash H \rightarrow \mathcal{W}'$ for some $H \in \vec{H}_j$ then $\Delta; \Phi \models \mathcal{W}'$.*

PROOF. The program $\text{prog} = \vec{F}_i \vec{H}_j$ must have been well typed according to prog-WF, and so we have $\vdash \Delta$ and $\vdash \Phi$ along with derivations for each F and H :

$$\begin{array}{ccccccc} \Delta; \Phi \vdash F_0 & \dots & \Delta; \Phi \vdash F_i & \dots & \Delta; \Phi \vdash F_n \\ \Delta; \Phi \vdash H_0 & \dots & \Delta; \Phi \vdash H_j & \dots & \Delta; \Phi \vdash H_m \end{array}$$

Let an $H = p(X_k) :- \vec{P}_\ell \in \vec{H}_j$ be given. We know that $\Delta; \Phi \vdash H$ by H -CLAUSE, i.e.:

$$\begin{array}{ccccccc} \cdot \vdash P_0 \triangleright \Gamma_1 & \dots & \Gamma_\ell \vdash P_\ell \triangleright \Gamma_{\ell+1} & \dots & \Gamma_p \vdash P_p \triangleright \Gamma' \\ p \subseteq \vec{\tau}_k \in \Phi & & \Gamma' \vdash \vec{X}_k, \vec{\tau}_k \triangleright \Gamma' & & \end{array}$$

Let \mathcal{W} be given such that $\Delta; \Phi \models \mathcal{W}$. It must have been the case that we stepped by CLAUSE, and so:

$$\begin{array}{ccccccc} \cdot \vdash P_0 \rightarrow \theta_1 & \dots & \theta_i \vdash P_i \rightarrow \theta_{i+1} & \dots & \theta_n \vdash P_n \rightarrow \theta \\ \mathcal{W}' = \mathcal{W}[p \mapsto \mathcal{W}(p) \cup \theta(\vec{X}_k)] & & & & \end{array}$$

By Lemma E.4(2), we know that $\Gamma_i \models \theta_i$ and $\Gamma' \models \theta$. We have $\vec{X}_k \subseteq \text{dom}(\Gamma')$ by Lemma E.5 on $\Gamma' \vdash \vec{X}_k, \vec{\tau}_k \triangleright \Gamma'$, we can conclude that $\vec{X}_k \subseteq \text{dom}(\theta)$ and that $\Delta; \Phi; \cdot \vdash \theta(\vec{X}_k) : \tau_k$ by Lemma E.3 on $\Gamma' \vdash \vec{X}_k, \vec{\tau}_k \triangleright \Gamma'$.

To see that $\Delta; \Phi \models \mathcal{W}'$, we need to see that adding $\theta(\vec{X}_k)$ to $\mathcal{W}(p)$ is safe. We already knew that $p \subseteq \vec{\tau}_k \in \Phi$ and $p \in \text{dom}(\mathcal{W})$; we have $k = k$ immediately, and we have seen that each $\theta(X_k)$ is well typed at τ_k . \square

F MODEL-THEORETIC SEMANTICS

We have focused on the operational semantics of Formulog, as it helps us to reason about type safety. However, since we have kept Formulog close to Datalog, it is also possible to give a model-theoretic semantics to a Formulog program. First, all ML functions and expressions are desugared into Datalog rules; this translation is relatively straightforward, with the trickiest part being the translation of non-mutually exclusive patterns occurring in match expressions. For each primitive

operator, we introduce a (possibly infinite) EDB relation that defines that operator; for example, the addition operator $+$ is represented through a ternary relation $\text{add}(x, y, z)$, which states that z is the sum of x and y . Terms of the form $p(w^*)$ (i.e., invocations of predicates as functions) are translated to aggregate predicates. Formulog requires the use of these terms, as well as negation, to be stratified; thus, the program resulting from the translation can be given a perfect model semantics in line with stratified negation [Apt et al. 1988; Przymusiński 1988; Van Gelder 1989] and stratified aggregation [Mumick et al. 1990].

For a small example, consider this Formulog program:

```
fun length(Xs: 'a list) : bv[32] =
  match Xs with
  | [] => 0
  | _ :: T => 1 + length(T)
  end
ok :- length([1, 2, 3]) = 3.
```

This would be translated into a program like this:

```
length([], 0).
length(_ :: T, Z) :-
  length(T, L),
  add(1, L, Z).
ok :- length([1, 2, 3], 3).
```

Note that the rules defining the `length` predicate violate the range restriction, and in fact define an infinite relation. This does not pose a fundamental problem to the model theory. To make this program evaluable, we could rewrite the `length` predicate's definition (and its uses) via the magic set transformation; the resulting relations would meet the range restriction. While the magic set transformation can turn a stratified program into a non-stratified program, there are techniques to either restore stratification [Meskes and Noack 1993] or correctly evaluate the non-stratified program [Mumick et al. 1990].