

Semantic subtyping with an SMT solver

GAVIN M. BIERMAN and ANDREW D. GORDON

Microsoft Research, Cambridge, UK
(e-mail: gmb@microsoft.com, adg@microsoft.com)

CĂTĂLIN HRIȚCU

Saarland University and University of Pennsylvania, Philadelphia, PA, USA
(e-mail: catalin.hritcu@gmail.com)

DAVID LANGWORTHY

Microsoft Corporation, Redmond, WA
(e-mail: dlan@microsoft.com)

Abstract

We study a first-order functional language with the novel combination of the ideas of refinement type (the subset of a type to satisfy a Boolean expression) and type-test (a Boolean expression testing whether a value belongs to a type). Our core calculus can express a rich variety of typing idioms; for example, intersection, union, negation, singleton, nullable, variant, and algebraic types are all derivable. We formulate a semantics in which expressions denote terms, and types are interpreted as first-order logic formulas. Subtyping is defined as valid implication between the semantics of types. The formulas are interpreted in a specific model that we axiomatize using standard first-order theories. On this basis, we present a novel type-checking algorithm able to eliminate many dynamic tests and to detect many errors statically. The key idea is to rely on a Satisfiability Modulo Theories solver to compute subtyping efficiently. Moreover, using a satisfiability modulo theories solver allows us to show the uniqueness of normal forms for non-deterministic expressions, provide precise counterexamples when type-checking fails, detect empty types, and compute instances of types statically and at run-time.

1 Introduction

This paper studies first-order functional programming in the presence of both refinement types (types qualified by Boolean expressions) and type-tests (Boolean expressions testing whether a value belongs to a type). The novel combination of type-test and refinement types appears in a recent commercial functional language, code-named M (Microsoft Corporation 2009), whose types correspond to relational schemas, and whose expressions compile to SQL queries. Refinement types are used to express SQL table constraints within a type system, and type-tests are useful for processing relational data, for example, by discriminating dynamically between different forms of union types. Still, although useful and extremely expressive, the combination of type-test and refinement types is hard to type-check using conventional syntax-driven subtyping rules. The preliminary implementation of M

uses such subtyping rules and has difficulty with certain sound idioms (such as uses of singleton and union types). Hence, type safety is enforced by dynamic checks, or not at all.

This paper studies the problem of type-checking code that uses type-tests and refinements via a core calculus, named Dminor, whose syntax is a small subset of M, and which is expressive enough to encode all the essential features of the full M language. In the remainder of this section, we elaborate on the difficulties of type-checking Dminor (and hence M), and outline our solution, which is to use semantic subtyping rather than syntactic rules.

1.1 Programming with type-test and refinement

The core types of Dminor are structural types for scalars, unordered collections, and records. (Following the database orientation of M, we refer to records as *entities*.) We write $S <: T$ for the subtype relation, which means that every value of type S is also of type T .

Two central primitives of Dminor are the following:

- A *refinement type*, $(x : T \text{ where } e)$, consists of the values x of T satisfying the Boolean expression e .
- A *type-test expression*, $e \text{ in } T$, returns **true** or **false** depending on whether or not the value of e belongs to type T .

As we shall see, many types are derivable from these primitive constructs and their combination. For example, the singleton type $[v]$, which contains just the value v , is derived as the refinement type $(x : \text{Any} \text{ where } x == v)$, where **Any** is the type of all values. The union type $T \mid U$, which contains the values of T together with the values of U , is derived as $(x : \text{Any} \text{ where } (x \text{ in } T) \mid\mid (x \text{ in } U))$.

Here is a snippet from a typical Dminor (and M) program for processing a domain specific language, a language of while-programs. The type is a union of different sorts of statements, each of which is an entity with a *kind* field of singleton type. (The snippet relies on an omitted – but similar – recursive type of arithmetic expressions.)

```
type Statement =
  {kind:["assignment"]; var: Text; rhs: Expression;} |
  {kind:["while"]; test:Expression; body:Statement;} |
  {kind:["if"]; test:Expression; tt:Statement; ff:Statement;} |
  {kind:["seq"]; s1:Statement; s2:Statement;} |
  {kind:["skip"]};
```

In languages influenced by HOPE (Burstall *et al.*, 1980), such as ML and Haskell, we would use the built-in notion of algebraic type to represent such statements. But like many data formats, including relational databases, semi-structured data, S-expressions, and JavaScript Object Notation (JSON) (Crockford, 2006), the data structures of M and Dminor do not take as primitive the idea of data tagged with data constructors. Instead, we need to follow an idiom such as shown above, of taking the union of entity types that include explicit tags that are given distinct singleton types.

If y has type `Statement`, we may process such data as follows:

```
((y.kind == "assignment")? y.var : "NotAssign")
```

Intuitively, this code is type-safe because it checks the `kind` field before accessing the `var` field, which is only present for assignment statements. More precisely, to type-check the then-branch `y.var` at type `Text`, we have $y : \text{Statement}$ (that is, a union type encoded using refinements and type-test, which after expansion has the form $(x : \text{Any} \textbf{where} \dots \parallel \dots)$), know that `y.kind == "assignment"`, and need to decide $[y] <: \{\text{var} : \text{Text};\}$. Subtyping should succeed, but clearly requires relatively sophisticated symbolic computation, including case analysis and propagation of equations. This is a typical example where syntax-driven rules for refinements and type-test are inadequate (if one ignores the refinement in the definition of `Statement`, then `Any` is not a subtype of $\{\text{var} : \text{Text};\}$), and indeed this simple example cannot be checked statically by the preliminary release of M. Our proposal is to delegate the hard work to an external prover.

1.2 An opportunity: SMT as a platform

Over the past few years, there has been tremendous progress in the field of Satisfiability Modulo Theories (SMT), that is, for (fragments of) first-order logic (FOL) plus various standard theories such as equality, real and integer (linear) arithmetic, bit vectors, and (extensional) arrays. Some of the leading systems include CVC3 (Barrett & Tinelli, 2007), Yices (Dutertre & de Moura, 2006), and Z3 (de Moura & Bjørner, 2008). There are common input formats such as Simplify's (Detlefs *et al.*, 2005) unsorted S-expression syntax and the SMT-LIB standard (Ranise & Tinelli, 2006) for sorted logic. Hence, FOL with standard theories is emerging as a computing platform. Software written to generate problems in a standard format can rely on a wide range of back-end solvers, which get better over time due in part to healthy competition,¹ and which may even be run in parallel when sufficient cores are available. There are limitations, of course, as first-order validity is undecidable even without any theories, so solvers may fail to terminate within a reasonable time, but recent progress has been remarkable.

1.3 Semantic subtyping with an SMT solver

The central idea in this paper is a type-checking algorithm for Dminor that checks subtyping by invoking an external SMT solver. To check whether S is a subtype of T , we construct first-order formulas $\mathbf{F}[\![S]\!](x)$ and $\mathbf{F}[\![T]\!](x)$, which hold when x belongs to the type S and the type T , respectively, and ask the solver whether the formula $\mathbf{F}[\![S]\!](x) \Rightarrow \mathbf{F}[\![T]\!](x)$ is valid, given any additional constraints known from the typing environment. This technique is known as *semantic subtyping* (Aiken & Wimmers, 1993; Frisch *et al.*, 2008), as opposed to the more common alternative, *syntactic*

¹ Most important is the SMT-COMP (Barrett *et al.*, 2008) competition held each year in conjunction with CAV and in which more than a dozen SMT solvers contend.

subtyping, which is to define syntax-driven rules for checking subtyping (Pierce, 2002).

The idea of using an external solver for type-checking with refinement types is not new. Several recent type-checkers for functional languages, such as SAGE (Flanagan, 2006; Knowles *et al.*, 2007; Knowles & Flanagan, 2010), F7 (Bengtson *et al.*, 2008), Fine (Swamy *et al.*, 2010), and Dsolve (Rondon *et al.*, 2008), rely on various SMT solvers. However, these systems all rely on syntactic subtyping, with the solver being used as a subroutine to check constraints during subtyping.

To the best of our knowledge, our proposal to implement semantic subtyping by calling an external SMT solver is new. Semantic subtyping nicely exploits the solver's ability to handle logical connectives efficiently; for example, we represent union and intersection types as logical disjunctions and conjunctions. Hence, we avoid the implementation effort of explicit propagation of constraints, and of syntax-driven rules for union and intersection types (Pierce, 1991; Dunfield & Pfenning, 2004; Dunfield, 2007). Moreover, we exploit the theories of equality, integer arithmetic, extensional arrays (de Moura & Bjørner, 2009), and algebraic datatypes.

1.4 Contributions of the paper

1. Investigation of semantic subtyping for a core functional language with both refinement types and type-test expressions (a novel combination, as far as we know). We are surprised that so many typing constructs are derivable from this combination.
2. Development of the theory, including both a declarative type assignment relation, and algorithmic rules in the bidirectional style. Our correctness results cover the core type assignment relation, the bidirectional rules, the algorithmic purity check, and some logical optimizations.
3. An implementation based on checking semantic subtyping by constructing proof obligations for an external SMT solver. The proof obligations are interpreted in a model that is formalized in Coq and axiomatized using standard first-order theories (equality, integers, datatypes, and extensional arrays).
4. Devising a systematic way to use the SMT solver in order to show the uniqueness of normal forms for non-deterministic expressions, provide precise counterexamples when type-checking fails, detect empty types, and compute instances of types. The latter enables a new form of declarative constraint programming, where constraints arise from the interpretation of a type as a formula.

1.5 Structure of the paper

Section 2 describes the formal syntax of Dminor together with a small-step operational semantics, $e \rightarrow e'$, where e and e' are expressions. We encode a series of type idioms to illustrate the expressiveness of the language and its type system.

Section 3 presents a logical semantics of pure expressions (those without side-effects, including non-termination) and Dminor types. We require that expressions used as refinements be pure so that they have a direct interpretation as predicates. Each pure expression e is interpreted as a term $\mathbf{R}[[e]]$ and each type T is interpreted as a FOL formula $\mathbf{F}[[T]](t)$, where t is a FOL term. The formulas are interpreted in a specific model that we have formalized in Coq. Theorem 1 is a full abstraction result: Two pure expressions have the same logical semantics just when they are operationally equivalent.

Section 4 introduces a tractable property, *algorithmic purity*, for use in our typing rules. Algorithmic purity is defined using a syntactic termination restriction together with a confluence check that relies on the logical semantics. Theorem 2 shows that our algorithmic purity check is indeed a sufficient condition for purity.

Section 5 presents the declarative type system for Dminor. The type assignment relation has the form $E \vdash e : T$, meaning that expression e has type T given typing environment E . Theorem 3 concerns logical soundness of type assignment; if e is assigned type T then formula $\mathbf{F}[[T]](\mathbf{R}[[e]])$ holds. Progress and preservation results (Theorems 4 and 5) relate type assignment to the operational semantics, entailing that well-typed expressions cannot go wrong.

Section 6 develops additional theory to justify our implementation techniques. First, we present simpler variations of the translations $\mathbf{R}[[e]]$ and $\mathbf{F}[[T]](t)$, optimized by the observation that during type-checking we only interpret well-typed expressions, and so we need not track error values. Theorem 6 shows the soundness and completeness of this optimization. Second, since the declarative rules of Section 5 are not directly algorithmic, we propose type-checking and synthesis algorithms presented as bidirectional rules. Theorem 7 shows these are sound with respect to type assignment.

Section 7 shows how to use the models produced by the SMT solver to provide very precise counterexamples when type-checking fails and to find inhabitants of types statically or dynamically. Section 8 reports some details of our implementation. We survey related work in Section 9, before concluding in Section 10.

The appendixes describe our intended logical model of Dminor and its formalization in Coq (Appendix A); report on the axiomatization of the model passed to the SMT solver during type-checking (Appendix B); and provide detailed proofs (Appendix C).

Our implementation, as well as sample code and listings of Dminor runs, the Coq formalization of our model, and also a screencast comparing the effectiveness of Dminor with the standard M type-checker are all available at <http://research.microsoft.com/dminor/>. A technical report (Bierman *et al.*, 2010a) describes how our type-checker may be used to check for systems configuration errors. A preliminary abridged version of this work appears in a conference proceedings (Bierman *et al.*, 2010b). Although we have formalized our logical model, and some other definitions, in Coq, our proofs are not in general mechanized in Coq. We list the theorems proved in Coq in Section 8.

Finally, we report that the future of the M language, the inspiration for Dminor, is rather uncertain at present. In a September 2010 blog posting (Box, 2010), Microsoft

announced that prototype software based on M would not be brought to market. Whatever the future of the M language itself, our hope is that the concepts we have developed in Dminor will be valuable in other settings.

2 Syntax and operational semantics

Dminor is a strict first-order functional language whose data include scalars, entities, and collections; it has no mutable state, and its only side effects are non-termination and non-determinism. This section describes (1) the syntax of expressions, types, and global function definitions; (2) the operational semantics; (3) the definition of pure expressions (those without side effects); and (4) some encodings to justify our expressiveness claims.

The following example introduces the basic syntax of Dminor. An accumulate expression is a fold over an unordered collection; to evaluate `from x in e1 let y = e2 accumulate e3`, we first evaluate `e1` to a collection `v`, evaluate `e2` to an initial value `u0`, and then compute a series of values `ui` for $i \in 1..n$ by setting `ui` to the value of `e3{vi/x}{ui-1/y}`, and eventually return `un`, where `v1, ..., vn` are the items in the collection `v`, in some arbitrary order.

`NullableInt` \triangleq `Integer` | [`null`]

`removeNulls(xs : NullableInt*) : Integer*`

`{ from x in xs let a = ({}:Integer*) accumulate (x!=null) ? (x :: a) : a }`

The type `NullableInt` is defined as the union of `Integer` with the singleton type containing only the value `null`. The type `Integer*` denotes a collection of values of type `Integer`. We then define a function `removeNulls` that iterates over its input collection and removes all `null` elements. As expected, executing `removeNulls({1, null, 42, null})` produces `{1, 42}` (which denotes the same collection as `{42, 1}`).

Given that the collection `xs` contains elements of type `NullableInt` (`xs : NullableInt*`), that `x` is an element of `xs`, and the check that `x != null`, our type-checking algorithm infers that on the if branch `x : Integer`, and therefore the result of the comprehension is `Integer*`, as declared by the function. If we remove the check that `x != null`, and copy all elements with `x :: a` then type-checking fails, as expected.

2.1 Expressions and types

We observe the following syntactic conventions. We identify all phrases of syntax (such as types and expressions) up to consistent renaming of bound variables. For any phrase of syntax ϕ we write $\phi\{e/x\}$ for the outcome of a capture-avoiding substitution of e for each free occurrence of x in ϕ . We write $fv(\phi)$ for the set of variables occurring free in ϕ .

We assume some base types for integers, strings, and logical values, together with constants for each of these types as well as a `null` value. We also assume an assortment of primitive operators; they are all binary apart from negation `!`, which is unary.

Scalar types, constants, and operators

$G ::= \text{Integer} \mid \text{Text} \mid \text{Logical}$	scalar type
$K(\text{Integer}) = \{i \mid \text{integer } i\}$	
$K(\text{Text}) = \{s \mid \text{string } s\}$	
$K(\text{Logical}) = \{\text{true}, \text{false}\}$	
$c \in K(\text{Integer}) \cup K(\text{Text}) \cup K(\text{Logical}) \cup \{\text{null}\}$	scalar constants
$\oplus \in \{+, -, \times, <, >, ==, !, \&\&, \}$	primitive operators

A *value* may be a *simple value* (an integer, string, Boolean, or **null**), a *collection* (a finite multiset of values), or an *entity* (a finite set of fields, each consisting of a value with a distinct label). (We follow M terminology, but entities would usually be called records.)

Syntax of values

$v ::=$	value
c	scalar (or simple value)
$\{v_1, \dots, v_n\}$	collection (multiset; unordered)
$\{\ell_i \Rightarrow v_i \mid i \in 1..n\}$	entity (ℓ_i distinct)

We identify values u and v , and write $u = v$ when they are identical up to reordering the items within collections or entities. While collections are unordered, ordered lists can be encoded using nested entities (see Section 2.4).

Syntax of types

$S, T, U ::=$	type
Any	the top type
G	scalar type
T^*	collection type
$\{\ell : T\}$	(single) entity type
$(x : T \text{ where } e)$	refinement type (scope of x is e)

All values have type **Any**, the top type. The values of a scalar type G are the scalars in the set $K(G)$ defined above. The values of type T^* are collections of values of type T . The values of type $\{\ell : T\}$ are entities with (at least) a field ℓ holding values of type T . (We show in Section 2.4 how to define multi-field entity types as a form of intersection type.) Finally, the values of a *refinement type* $(x : T \text{ where } e)$ are the values v of type T such that the Boolean expression $e\{v/x\}$ returns **true**. As a convenient shorthand, we write $T \text{ where } e$ for the refinement type (value : $T \text{ where } e$), where the omitted variable defaults to **value**. For example, **Integer where value > 0** is the type of positive numbers.

Syntax of expressions

$e ::=$	expression
x	variable
c	scalar constant

Derived LINQ queries

from x **in** e_1 **where** e_2 **select** $e_3 \triangleq$ **from** x **in** e_1 **let** $y = \{\}$ **accumulate** $e_2?(e_3 :: y) : y$

To complete the syntax of Dminor, we interpret types and expressions in the context of a fixed collection of first-order, dependently typed, potentially recursive function definitions. We assume for each expression $f(e_1, \dots, e_n)$ in a source program that there is a corresponding function definition for f with arity n .

Function definitions: $f(x_1 : T_1, \dots, x_n : T_n) : U\{e\}$

We assume a finite, global set of *function definitions*, each of which associates a function name f with a dependent signature $x_1 : T_1, \dots, x_n : T_n \rightarrow U$, formal parameters x_1, \dots, x_n , and a body e , such that $fv(e) \subseteq \{x_1, \dots, x_n\}$ and $fv(U) \subseteq \{x_1, \dots, x_n\}$.

2.2 Operational semantics

We define a non-deterministic, potentially divergent, small-step reduction relation $e \rightarrow e'$, together with a standard notion of expressions going wrong, to be prevented by typing.

Each primitive operator is a partial function represented by a set of mappings of the form $\oplus(v_1, \dots, v_n) \mapsto v_0$, where each v_i is a value. The $==$ operator implements syntactic equality, which for collections and entities is up to reordering of elements. Apart from $==$, the other operators only act on scalar values. For example, the equations for $+$ are $(\underline{i} + \underline{j}) \mapsto \underline{i + j}$. The other operators are defined by similar equations, and we omit the details.

Reduction contexts

$\mathcal{R} ::=$ reduction context
 $\oplus(v_1, \dots, v_{j-1}, \bullet, e_{i+1}, \dots, e_n)$
 $\bullet?e_2 : e_3 \mid \mathbf{let} \ x = \bullet \ \mathbf{in} \ e_2 \mid \bullet \ \mathbf{in} \ T$
 $\{\ell_i \Rightarrow v_i \ i \in 1..j-1, \ell_j \Rightarrow \bullet, \ell_i \Rightarrow e_i \ i \in j+1..n\}$
 $\bullet.\ell \mid \bullet :: e \mid v :: \bullet \mid \mathbf{from} \ x \ \mathbf{in} \ \bullet \ \mathbf{let} \ y = e_2 \ \mathbf{accumulate} \ e_3$
 $f(v_1, \dots, v_{i-1}, \bullet, e_{i+1}, \dots, e_n)$

Reduction rules for standard constructs

$e \rightarrow e' \Rightarrow \mathcal{R}[e] \rightarrow \mathcal{R}[e']$
 $\oplus(v_1, \dots, v_n) \rightarrow v$ if $\oplus(v_1, \dots, v_n) \mapsto v$ defined
 $\mathbf{true}?e_2 : e_3 \rightarrow e_2$
 $\mathbf{false}?e_2 : e_3 \rightarrow e_3$
 $\mathbf{let} \ x = v \ \mathbf{in} \ e_2 \rightarrow e_2\{v/x\}$
 $\{\ell_i \Rightarrow v_i \ i \in 1..n\}.\ell_j \rightarrow v_j$ where $j \in 1..n$
 $v :: \{v_1, \dots, v_n\} \rightarrow \{v_1, \dots, v_n, v\}$
 $\mathbf{from} \ x \ \mathbf{in} \ \{v_1, \dots, v_n\} \ \mathbf{let} \ y = e_2 \ \mathbf{accumulate} \ e_3$
 $\rightarrow \mathbf{let} \ y = e_2 \ \mathbf{in} \ \mathbf{let} \ y = e_3\{v_1/x\} \ \mathbf{in} \ \dots \ \mathbf{let} \ y = e_3\{v_n/x\} \ \mathbf{in} \ y$

$$f(v_1, \dots, v_n) \rightarrow e\{v_1/x_1\} \dots \{v_n/x_n\}$$

given function definition $f(x_1 : T_1, \dots, x_n : T_n) : U\{e\}$

Reduction rules for type-test

$$v \text{ in Any} \rightarrow \mathbf{true}$$

$$v \text{ in } G \rightarrow \begin{cases} \mathbf{true} & \text{if } v \in K(G) \\ \mathbf{false} & \text{otherwise} \end{cases}$$

$$v \text{ in } \{\ell_j : T_j\} \rightarrow \begin{cases} v_j \text{ in } T_j & \text{if } v = \{\ell_i \Rightarrow v_i \mid i \in 1..n\} \wedge j \in 1..n \\ \mathbf{false} & \text{otherwise} \end{cases}$$

$$v \text{ in } T^* \rightarrow \begin{cases} v_1 \text{ in } T \ \&\& \dots \ \&\& v_n \text{ in } T & \text{if } v = \{v_1, \dots, v_n\} \\ \mathbf{false} & \text{otherwise} \end{cases}$$

$$v \text{ in } (x : T \ \mathbf{where} \ e) \rightarrow v \text{ in } T \ \&\& \ e\{v/x\}$$

The reduction rules for type-test expressions, $e \text{ in } U$, first reduce e to a value v and then proceed by case analysis on the structure of the type U . In case U is a refinement type $(x : T \ \mathbf{where} \ e)$ then v is a value of U if and only if v is a value of type T and $e\{v/x\}$ reduces to the value \mathbf{true} .

The reduction relation would be deterministic were it not for the reduction rule for accumulate expressions. (If the primitive syntax for collections was not $\{v_1, \dots, v_n\}$ but instead was $\{e_1, \dots, e_n\}$, where the e_i are not necessarily values, non-determinism would also arise from the reduction of collections to values.) Since collections are unordered, the rule applies for any permutation of $\{v_1, \dots, v_n\}$. For example, consider the expression $\text{pick } v_1 \ v_2 \triangleq \mathbf{from } x \text{ in } \{v_1, v_2\} \ \mathbf{let } y = \mathbf{null} \ \mathbf{accumulate } x$; we have both $\text{pick } \mathbf{true} \ \mathbf{false} \rightarrow^* \mathbf{true}$ and $\text{pick } \mathbf{true} \ \mathbf{false} \rightarrow^* \mathbf{false}$.

Next, we use reduction to define an evaluation relation that relates a closed expression to its return values, or to \mathbf{Error} , in case reduction gets stuck before reaching a value.

Stuckness, results, and evaluation: $e \Downarrow r$ for closed e

Let e be *stuck* if and only if e is not a value and $\neg \exists e'. e \rightarrow e'$.

$r ::= \mathbf{Error} \mid \mathbf{Return}(v)$ results of evaluation

$e \Downarrow \mathbf{Return}(v)$ if and only if $e \rightarrow^* v$

$e \Downarrow \mathbf{Error}$ if and only if there is e' such that $e \rightarrow^* e'$ and e' is stuck.

Let closed expression e *go wrong* if and only if $e \Downarrow \mathbf{Error}$. For example, we have that $\text{stuck} \Downarrow \mathbf{Error}$, where $\text{stuck} \triangleq \{\}. \ell$ for some label ℓ . In the presence of type-test and refinement types, expressions can go wrong in unusual ways. For example, given the refinement type $T = (x : \text{Any} \ \mathbf{where} \ \text{stuck})$, any type-test $v \text{ in } T$ goes wrong. The main goal of our type system is to ensure that no closed well-typed expression goes wrong.

Encoding type-assertion

$$\mathbf{assert}(e : T) \triangleq \mathbf{let } x = e \ \mathbf{in} \ ((x \text{ in } T)?x : \text{stuck})$$

Using type-tests we can easily encode *type-assertions*. The expression `assert($e : T$)` enforces that the result of the expression e is a value of type T . Operationally, `assert($e : T$)` returns the value of e if this is an element of T , and goes wrong otherwise. Since our type system ensures that well-typed expressions do not go wrong, it also ensures statically that type-assertions always succeed.

Calling a function with arguments that do not have their declared types does *not* necessarily go wrong. Similarly, the operational semantics does not force functions to return a result that matches the declared type. One can, however, use explicit type-assertions to enforce that the declared types are respected by rewriting any function definition $f(x_1 : T_1, \dots, x_n : T_n) : U\{e\}$ into $f(x_1 : T_1, \dots, x_n : T_n) : U\{(x_1 \text{ in } T_1 \&\& \dots \&\& x_n \text{ in } T_n) ? \text{assert}(e : U) : \text{stuck}\}$. Our type system enforces in any case that declared types are respected. This enables us to express pre- and post-conditions of functions using refinement types.

2.3 Pure expressions and refinement types

A problem in languages with refinement types ($x : T$ **where** e) is that the refinement expression e , even though well-typed, may have effects, such as non-termination or non-determinism, and so makes no sense as a Boolean condition. In Dminor calls to recursive functions can cause divergence, and since collections are unordered, iterating over them with `accumulate` may be non-deterministic as above.

To address this problem, we define the set of *pure* expressions, the ones that may be used as refinements. The details below are little technical, but the gist is that pure expressions must be terminating, have a unique result (which may be **Error**), and must only call functions whose bodies are pure. The typing rule (Type Refine) in Section 5 requires that for ($x : T$ **where** e) to be well-formed, the expression e must be pure and of type Logical (which guarantees that e yields **true** or **false** without getting stuck). Checking for purity is undecidable, but we present sufficient conditions for checking purity algorithmically in Section 4.

We assume that a subset of the function definitions is *labeled-pure*; we intend that only these functions may be called from pure expressions. Let an expression e be *terminating* if and only if there exists no unbounded sequence $e \rightarrow e_1 \rightarrow e_2 \rightarrow \dots$. Let a closed expression e be *pure* if and only if (1) e is terminating, (2) there exists a unique result r such that $e \Downarrow r$, (3) for every subexpression $f(e_1, \dots, e_n)$ of e , the function f is labeled-pure, and (4) all subexpressions of e are pure. Let an arbitrary expression e be *pure* if and only if $e\sigma$ is pure for all closing substitutions σ that assign a value to each free variable in e . Finally, we require that the body of every labeled-pure function is a pure expression.

2.4 Derived types

We end this section by exploring the expressiveness of the primitive types introduced above, and in particular of the combination of refinement types and dynamic type-test. We show that the range of derivable types is rather wide. We begin with some basic examples.

Encoding of empty and singleton types

$$\text{Empty} \triangleq (x : \text{Any} \text{ where } \text{false})$$

$$[e] \triangleq (x : \text{Any} \text{ where } x == e) \quad (e \text{ pure}, x \notin \text{fv}(e))$$

The type `Empty` has no elements; it is a subtype of all other types. The *singleton type*, `[e]`, contains only the value of pure expression e (for example, type `[null]` consists just of the `null` value).

Our calculus includes the operators of propositional logic on Boolean values. We lift these operators to act on types as follows.

Encoding of union, intersection, and negation types

$$T \mid U \triangleq (x : \text{Any} \text{ where } (x \text{ in } T) \mid (x \text{ in } U)) \quad (x \notin \text{fv}(T, U))$$

$$T \ \& \ U \triangleq (x : \text{Any} \text{ where } (x \text{ in } T) \ \&\& \ (x \text{ in } U))$$

$$!T \triangleq (x : \text{Any} \text{ where } !(x \text{ in } T))$$

A value of the *union type*, $T \mid U$, is a value of T or of U . A value of the *intersection type*, $T \ \& \ U$, is a value of both T and U . A value of the *negation type*, $!T$, is a value that is not a value of T . We omit the details, but we could go in the other direction too: Boolean operators are derivable from union, intersection, and complement types.

Next, we define the types of simple values, collections, and entities. We rely on the primitive types `Integer`, `Text`, and `Logical`, the primitive type constructor T^* for collections, and the fact that every proper value is either a scalar, a collection, or an entity: so the type of entities is the complement of the union type `General` \mid `Collection`.

Encoding of supertypes

$$\text{General} \triangleq \text{Integer} \mid \text{Text} \mid \text{Logical} \mid [\text{null}]$$

$$\text{Collection} \triangleq \text{Any}^*$$

$$\text{Entity} \triangleq !(\text{General} \mid \text{Collection})$$

(This encoding illustrates the power of types based on propositional logic, but is fragile; if we were to extend the language with other primitive types, it would be better to take `Entity` to be primitive too, rather than defining it as a complement.)

The primitive type of entities is unary: the type $\{\ell : T\}$ is the set of entities with a field ℓ whose value belongs to T (and possibly other fields). As in Forsythe (Reynolds, 1996), we derive *multiple-field entity types* as an intersection type. One advantage of this approach is that it immediately entails width and depth subtyping for entities.

Encoding of multiple-field entity types

$$\{\ell_i : T_i; i \in 1..n\} \triangleq \{\ell_1 : T_1\} \ \& \ \dots \ \& \ \{\ell_n : T_n\} \quad (\ell_i \text{ distinct}, n > 0)$$

We can also derive *closed entity types*, which only contain entities with a fixed set of labels, and therefore allow depth but not width subtyping. To do this we constrain the multiple-field entity types above to additionally satisfy an eta law.

Encoding of closed entity types

$$\mathbf{closed}\{\ell_i : T_i; \ i \in 1..n\} \triangleq (x : \{\ell_i : T_i; \ i \in 1..n\} \mathbf{where} \ x == \{\ell_i \Rightarrow x.\ell_i \ i \in 1..n\})$$

Pair types are just a special case of closed entity types. Given pair types, refinement types, and type-test, we can also encode *dependent pair types* $\Sigma x : T. U$, where x is bound in U .

Encoding of pair types and dependent pair types

$$T * U \triangleq \mathbf{closed}\{\mathbf{fst} : T; \mathbf{snd} : U;\}$$

$$(\Sigma x : T. U) \triangleq (p : T * \mathbf{Any} \mathbf{where} \ \mathbf{let} \ x = p.\mathbf{fst} \ \mathbf{in} \ (p.\mathbf{snd} \ \mathbf{in} \ U))$$

Sum types are obtained from union types by adding an additional Boolean tag; *variant types* are a generalization.

Encoding of sum and variant types

$$T + U \triangleq ([\mathbf{true}] * T) \mid ([\mathbf{false}] * U)$$

$$\langle \ell_1 : T_1; \dots; \ell_n : T_n \rangle \triangleq \mathbf{closed}\{\ell_1 : T_1\} \mid \dots \mid \mathbf{closed}\{\ell_n : T_n\}$$

Recursive types can be encoded as Boolean recursive functions that dynamically test whether a given value has the required type.

Encoding recursive types

$$\mu X.T \triangleq (x : \mathbf{Any} \mathbf{where} \ f_{\mu X.T}(x)), \text{ where } f_{\nu}(T) = \emptyset$$

and $f_{\mu X.T}(x)$ is a new labeled-pure function defined by

$$f_{\mu X.T}(x : \mathbf{Any}) : \mathbf{Logical} \ \{x \ \mathbf{in} \ T\{(x : \mathbf{Any} \ \mathbf{where} \ f_{\mu X.T}(x))/X\}\}$$

The usual contractivity condition is replaced by the requirement that $f_{\mu X.T}$ is labeled-pure.

Using recursive, sum, and pair types we can encode any *algebraic datatype*. For instance, the type of lists of elements of type T can be encoded as follows.

Encoding list types

$$\mathbf{List}_T \triangleq \mu X. (T * X) + [\mathbf{null}]$$

In this encoding of lists, the cons-cell $v :: u$ is represented by a couple of nested entities $\{\mathbf{fst} \Rightarrow \mathbf{true}, \mathbf{snd} \Rightarrow \{\mathbf{fst} \Rightarrow v, \mathbf{snd} \Rightarrow u\}\}$. More efficient representations can be easily supported, for instance, $\mu X. \mathbf{closed}\{\mathbf{hd} : T, \mathbf{tl} : X\} \mid [\mathbf{null}]$, for which a cons-cell is a single entity: $\{\mathbf{hd} \Rightarrow v, \mathbf{tl} \Rightarrow u\}$.

Lists can be used to encode XML and JSON. Hence, Dminor can be viewed as a richly typed functional notation for manipulating data in XML format. And while, Document Type Definitions can be encoded as Dminor types, XML data can be loaded into Dminor even if there is no prior schema. We map an XML element to an entity, with a field to represent the name of the element, additional fields for any attributes on the element, and a final field holding a list of all the items in the body of the element.

Next, we show how to derive entity types for the common situation where the type of one field depends on the value of another. A *dependent intersection type* $(x : T \ \& \ U)$ (Kopylov, 2003) is essentially the intersection of T and U , except that the variable x is bound to the underlying value, with scope U . The type T cannot mention x , but we can rely on $x : T$ when checking well-formedness of U .

Encoding of dependent intersection types

$$(x : T \ \& \ U) \triangleq (x : T \ \mathbf{where} \ x \ \mathbf{in} \ U)$$

With this construct, we can define entity types where the type of one field depends on the value of another. For example, $(p : \{X : \text{Integer}\} \ \& \ \{Y : (y : \text{Integer} \ \mathbf{where} \ y < p.X)\})$ is the type of points below the diagonal.

M allows the field names of previously defined fields to be used within the types of subsequent fields. We can encode M's dependent entity types as follows.

Encoding dependent entities

$$\{\ell : T; \}U \triangleq (x : \{\ell : T\} \ \& \ U\{x.\ell/\ell\}) \quad \mathbf{where} \ x \notin \mathit{fv}(T, U)$$

$$\{\ell_1 : T_1; \dots; \ell_n : T_n; \} \triangleq \{\ell_1 : T_1; \} \dots \{\ell_n : T_n; \} \mathbf{Any}$$

Our example type of points below the diagonal is written in M as follows (where the field name X appears directly as an expression in the type of the field Y).

$$\{X : \text{Integer}; Y : \text{Integer} \ \mathbf{where} \ Y > X; \}$$

Our encoding turns this M syntax into the following type, which is equivalent to the more direct encoding given above.

$$(x_1 : \{X : \text{Integer}\} \ \& \ (x_2 : \{Y : \text{Integer} \ \mathbf{where} \ Y > x_1.X\} \ \& \ \mathbf{Any}))$$

Kopylov, (2003) explains in detail the relationship between dependent intersection and encodings of dependent entities (records).

To further illustrate the power of collection types combined with refinements, we give types below that express universal and existential quantifications over the items in a collection. Collection $\{v_1, \dots, v_n\} : T^*$ has type $\mathit{all}(x : T)e$ if $e\{v_i/x\}$ for all $i \in 1..n$, and, dually, it has type $\mathit{exists}(x : T)e$ if $e\{v_i/x\}$ for some $i \in 1..n$.

Quantifying over collections

$$\mathit{all}(x : T)e \triangleq (x : T \ \mathbf{where} \ e)^* \quad (e \ \mathbf{pure})$$

$$\mathit{exists}(x : T)e \triangleq T^* \ \& \ \mathit{!(all}(x : T)\mathit{!}e)$$

Curiously, a Boolean test for whether a value is a member of a collection need not be primitive in the calculus; we can make use of the type $\mathit{exists}(x : \mathbf{Any})(x == e_i)$ of collections that contain the item e_i , as follows.

Collection membership as a type-test

$$\mathit{Mem}(e_i, e_c) \triangleq (e_c \ \mathbf{in} \ \mathit{exists}(x : \mathbf{Any})(x == e_i)) \quad (e_i \ \mathbf{pure})$$

The Boolean expression $\text{Mem}(e_i, e_c)$ holds just when the value of e_i is a member of the collection denoted by e_c . (This example is to illustrate the expressiveness of the type system; collection membership is definable more directly by using an **accumulate** expression as shown in Section 2.1.)

The following dependent entity type consists of a collection of song titles `Songs`, together with a default. The type includes the constraint that the default song is a member of `Songs`.

```
{Songs: Text*; Default: Text where Mem(value,Songs);}
```

3 Logical semantics

In this section we give a set-theoretic semantics for types and pure expressions. Pure expressions are interpreted as first-order terms, while types are interpreted as formulas in many-sorted FOL. These formulas are interpreted in a fixed model, which we formalize in Coq. We represent a Dminor subtyping problem as a logical implication, supply our SMT solver with a set of axioms that are true in our intended model, and ask the solver to prove the validity of the implication. We use Coq to state our model and to reason about the soundness of the axioms given to the SMT solver, but semantic subtyping calls only the SMT solver, not Coq.

To represent the intended logical model formally, sets are encoded as Coq types, and functions are encoded as Coq functions. We start by encoding scalars, values and results, which were given as grammars in Section 2, as Coq types `General`, `Value` and `Result`.

Model: scalars, values, and results

Inductive `General` : Type :=

```
| G_Integer : Z → General
| G_Text : string → General
| G_Logical : bool → General
| G_Null : General.
```

Inductive `RawValue` : Type :=

```
| G : General → RawValue
| E : list (string * RawValue) → RawValue
| C : list RawValue → RawValue.
```

Definition `Value` := {x : RawValue | Normal x}.

Inductive `Result` : Type :=

```
| Error : Result
| Return : Value → Result.
```

Scalars and results are represented directly as Coq inductive types, while for values additional care is needed to prevent duplicate labels in entities and to ensure that the representation is canonical. Our Coq representation of values is explained in detail in Appendix A (for example, we define `Normal` in Appendix A), but the precise details are not relevant at this point. For the sake of readability, in this section we continue to use the intuitive notation for values and results introduced in Section 2.

We define a predicate `Proper` that is true for results that are not `Error`, and a function `out_V` that returns the value inside if the result passed as argument is `proper`

and **null** otherwise (the functions in the model are total, so in cases like this we return an arbitrary value).

Model: proper results

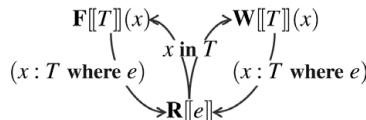
Definition Proper (res : Result) :=
match res **with** | **Return** v => **true** | **Error** => **false** **end**.
Definition out_V (res : Result) : Value :=
match res **with** | **Return** v => v | **Error** => v_null **end**.

Our semantics uses many-sorted FOL, and each sort is interpreted by a Coq type of the same name. We write predicates as functions to sort bool (interpreted by type bool in Coq), with truth values **true** and **false** – we let the context disambiguate between the truth values in the model and the corresponding Dminor Boolean values. We assume a collection of sorted function symbols whose interpretation in the intended model is given below. Let t range over FOL terms; we write $t : \sigma$ to mean that term t has sort σ ; if we omit the sort of a bound variable, it may be assumed to be Value. Similarly, free variables have sort Value by default. If F is a formula, let $\models F$ means that F is valid in our intended model.

Our semantics consists of three translations:

- For any expression e that only calls labeled-pure functions, we have the FOL term $\mathbf{R}[[e]] : \text{Result}$.
- For any Dminor type T and FOL term $t : \text{Value}$, we have the FOL formula $\mathbf{F}[[T]](t)$, which is valid in the intended model if and only if the value denoted by t is a member of the type T .
- For type T and FOL term $t : \text{Value}$, we have the formula $\mathbf{W}[[T]](t)$, which holds if and only if a type-test goes wrong when checking whether the value denoted by t is a member of T . For instance, we have $\models \mathbf{W}[(x : \text{Any where stuck})](\text{null}) \Leftrightarrow \text{true}$, but $\models \mathbf{W}[[\text{Any}]](\text{null}) \Leftrightarrow \text{false}$.

The relations between translations



These three (mutually recursive) translations are defined below. We rely on notations for let-binding within terms (**let** $x = t$ **in** t'), and terms conditional on formulas (**if** F **then** t **else** t'). These notations are supported directly by most SMT solvers. They can be translated to pure FOL by introducing auxiliary definitions, but we omit the details. Given these we can define the monadic bind for propagating

errors as a simple notation. Notice that $\models (\text{Bind } x \leftarrow \text{Return}(v) \text{ in } t) = t\{v/x\}$ and $\models (\text{Bind } x \leftarrow \text{Error in } t) = \text{Error}$.

Notation: monadic bind for propagating errors

$$\text{Bind } x \leftarrow t_1 \text{ in } t_2 \triangleq (\text{if } \neg \text{Proper}(t_1) \text{ then Error else let } x = \text{out.V}(t_1) \text{ in } t_2)$$

We begin by describing the semantics of some core types and expressions. The semantics of refinement types $\mathbf{F}[(x : T \text{ where } e)](t)$ relies on the result of evaluating e with x bound to t . Remember, however, that operationally the type test $v \text{ in } (x : T \text{ where } e)$ evaluates to **Error** if $e\{v/x\}$ evaluates to **Error** or to a value that is not **true** or **false**. We use $\mathbf{W}[(x : T \text{ where } e)](t)$ to record this fact, and we enforce that $\mathbf{R}[e \text{ in } T]$ returns **Error** if $\mathbf{W}[T](t)$ holds. Tracking type-tests going wrong is crucial for our full-abstraction result.

Semantics: core types and expressions

$$\mathbf{F}[\text{Any}](t) = \text{true}$$

$$\mathbf{W}[\text{Any}](t) = \text{false}$$

$$\mathbf{F}[(x : T \text{ where } e)](t) = \mathbf{F}[T](t) \wedge \text{let } x = t \text{ in } (\mathbf{R}[e] = \text{Return}(\text{true}))$$

$$\mathbf{W}[(x : T \text{ where } e)](t) = \mathbf{W}[T](t) \vee$$

$$\text{let } x = t \text{ in } (\neg(\mathbf{R}[e] = \text{Return}(\text{false})) \vee \mathbf{R}[e] = \text{Return}(\text{true}))$$

$$\mathbf{R}[x] = \text{Return}(x)$$

$$\mathbf{R}[e_1 ? e_2 : e_3] = \text{Bind } x \leftarrow \mathbf{R}[e_1] \text{ in}$$

$$(\text{if } x = \text{true} \text{ then } \mathbf{R}[e_2] \text{ else } (\text{if } x = \text{false} \text{ then } \mathbf{R}[e_3] \text{ else Error}))$$

$$\mathbf{R}[\text{let } x = e_1 \text{ in } e_2] = \text{Bind } x \leftarrow \mathbf{R}[e_1] \text{ in } \mathbf{R}[e_2]$$

$$\mathbf{R}[e \text{ in } T] = \text{Bind } x \leftarrow \mathbf{R}[e] \text{ in } (\text{if } \mathbf{W}[T](x) \text{ then Error else}$$

$$(\text{if } \mathbf{F}[T](x) \text{ then Return}(\text{true}) \text{ else Return}(\text{false})))$$

Next, we specify the semantics of scalar types and values. Function `is_G` in the model tests whether a value is a scalar or not, and if this is the case `out_G` returns this scalar. Similarly, functions `is_G_Logical`, `is_G_Integer`, and `is_G_Text` test whether a scalar has the corresponding scalar type.

Model: testers for simple values

Definition `In_Logical v := (is_G v) && is_G_Logical (out_G v).`

Definition `In_Integer v := (is_G v) && is_G_Integer (out_G v).`

Definition `In_Text v := (is_G v) && is_G_Text (out_G v).`

Semantics: scalar types, simple values, and operators

$$\mathbf{F}[\text{Integer}](t) = \text{In_Integer}(t)$$

$$\mathbf{R}[c] = \text{Return}(c)$$

$$\mathbf{F}[\text{Text}](t) = \text{In_Text}(t)$$

$$\mathbf{W}[G](t) = \text{false}$$

$$\mathbf{F}[\text{Logical}](t) = \text{In_Logical}(t)$$

$$\mathbf{R}[\oplus(e_1, \dots, e_n)] = \text{Bind } x_1 \leftarrow \mathbf{R}[e_1] \text{ in } \dots \text{Bind } x_n \leftarrow \mathbf{R}[e_n] \text{ in}$$

$$(\text{if } \mathbf{F}[T_1](x_1) \wedge \dots \wedge \mathbf{F}[T_n](x_n) \text{ then Return}(\text{O}_{\oplus}(x_1, \dots, x_n)) \text{ else Error})$$

where $\oplus : T_1, \dots, T_n \rightarrow T$

The semantics of primitive operators on simple values is defined uniformly. We state below the signature $\oplus : T_1, \dots, T_n \rightarrow T$ for each operator \oplus . We also name a Coq function O_\oplus to define the meaning of each operator. Then we define the semantics $\mathbf{R}[\oplus(e_1, \dots, e_n)]$ of operator expressions. Each of the functions O_\oplus is defined in Appendix A.2.

Model: signatures ($\oplus : T_1, \dots, T_n \rightarrow T$) and semantics (O_\oplus)

$+$: Integer, Integer \rightarrow Integer	$O_+ = O_Add$
$-$: Integer, Integer \rightarrow Integer	$O_- = O_Minus$
\times : Integer, Integer \rightarrow Integer	$O_\times = O_Mult$
$<$: Integer, Integer \rightarrow Logical	$O_< = O_LT$
$>$: Integer, Integer \rightarrow Logical	$O_> = O_GT$
$==$: Any, Any \rightarrow Logical	$O_{==} = O_EQ$
$!$: Logical \rightarrow Logical	$O_! = O_Not$
$\&\&$: Logical, Logical \rightarrow Logical	$O_{\&\&} = O_And$
$ $: Logical, Logical \rightarrow Logical	$O_{ } = O_Or$

The semantics of an entity type $\{\ell : T\}$ is the set of all values (denoted by t) that are proper entities ($\text{Good_E}(t)$) having the field ℓ ($\text{v_has_field}(\ell, t)$), which contains a value of type T ($\mathbf{F}[[T]](\text{v_dot}(t, \ell))$). The model functions v_has_field and v_dot use the Coq library function `TheoryList.assoc` to obtain the value associated with a given key in a list of pairs. Similarly, a type-test of the form v in $\{\ell : T\}$ goes wrong only when v is an entity having a field ℓ that contains a value v_ℓ for which the type-test v_ℓ in T goes wrong. If v is not an entity having field ℓ then the type-test v in $\{\ell : T\}$ will simply return **false**, and not go wrong. This is reflected in the definition of $\mathbf{W}[\{\ell : T\}]$ below.

Model: functions and predicates on entities

Program Definition v_has_field (s : string) (v : Value) : bool :=
match `TheoryList.assoc eq_str_dec s` (out.E v) **with**
 | Some $v \Rightarrow$ **true** | None \Rightarrow **false** **end**.

Program Definition v_dot (s : string) (v : Value) : Value :=
match `TheoryList.assoc eq_str_dec s` (out.E v) **with**
 | Some $v \Rightarrow v$ | None \Rightarrow v_null **end**.

Semantics: entity types and expressions

$\mathbf{F}[\{\ell : T\}](t) = \text{Good_E}(t) \wedge \text{v_has_field}(\ell, t) \wedge \mathbf{F}[[T]](\text{v_dot}(t, \ell))$
 $\mathbf{W}[\{\ell : T\}](t) = \text{Good_E}(t) \wedge \text{v_has_field}(\ell, t) \wedge \mathbf{W}[[T]](\text{v_dot}(t, \ell))$
 $\mathbf{R}[\{\ell_i \Rightarrow e_i \text{ }^{i \in 1..n}\}] = \mathbf{Bind} \ x_1 \leftarrow \mathbf{in} \ \mathbf{R}[[e_1]] \dots \mathbf{Bind} \ x_n \leftarrow \mathbf{R}[[e_n]] \ \mathbf{in} \ \mathbf{Return}(\{\ell_i \Rightarrow x_i \text{ }^{i \in 1..n}\})$
 $\mathbf{R}[[e.\ell]] = \mathbf{Bind} \ x \leftarrow \mathbf{R}[[e]] \ \mathbf{in}$
 (if $\text{is_E}(x) \wedge \text{v_has_field}(\ell, x)$ then $\mathbf{Return}(\text{v_dot}(x, \ell))$ else **Error**)

The semantics of **from** x **in** e_1 **let** $y = e_2$ **accumulate** e_3 relies on a function `res_accumulate` that folds over a collection by applying a function of sort `ClosureRes2`, and if no error occurs at any step it returns a value, otherwise it returns **Error**. If the accumulate expression is pure, it produces the same result no matter what order

is used when folding. The model of the sort `ClosureRes2` is the set of functions from `Value` to `Value` to `Result`. We write the lambda-abstraction $\text{fun } x \ y \rightarrow \mathbf{R}[\![e_3]\!]$ for such a function. There are several standard techniques for representing lambda-abstractions in FOL (Meng & Paulson, 2008). Our implementation generates a fresh function symbol to represent each lambda-abstraction occurring in its input as a closure of sort `ClosureRes2`.

Model: functions and predicates on collections

Program Definition `v_mem` (`v cv : Value`) : `bool` :=
`TheoryList.mem eq_rval_dec v (out_C cv)`.

Program Definition `v_add` (`v cv : Value`) : `Value` :=
`C (insert_in_sorted_vb v (out_C cv))`.

Definition `ClosureRes2` := `Value` \rightarrow `Value` \rightarrow `Result`.

Program Fixpoint `res_acc_fold` (`f : ClosureRes2`) (`vb : VBag`) (`a : Result`) {measure `List.length vb`} : `Result` :=

```

match vb with
| nil  $\Rightarrow$  a
| v :: vb'  $\Rightarrow$  match a with Return va  $\Rightarrow$  res_acc_fold vb' (f va v)
| Error  $\Rightarrow$  Error end

```

end.

Definition `res_accumulate` (`f : ClosureRes2`) (`cv v : Value`) : `Result` :=
if `is_C cv` **then** `res_acc_fold f (out_C cv)` (**Return** v) **else** **Error**.

The semantics of the collection type T^* is the set of all values (denoted by t) that are proper collections (`Good_C(t)`) containing only elements of type T ($\forall x.v_mem(x, t) \Rightarrow \mathbf{F}[\![T]\!](x)$). On the other hand, a type-test goes wrong for a collection type T^* if the value being tested is a proper collection containing *some* value that causes evaluation to go wrong when testing whether it belongs to type T . The $\mathbf{F}[\![T^*]\!]$ and $\mathbf{W}[\![T^*]\!]$ cases are the only ones in our semantics that generate logical formulas containing first-order quantifiers.

Semantics: collection types and expressions

$\mathbf{F}[\![T^*]\!](t) = \text{Good_C}(t) \wedge (\forall x.v_mem(x, t) \Rightarrow \mathbf{F}[\![T]\!](x)) \quad x \notin fv(T, t)$

$\mathbf{W}[\![T^*]\!](t) = \text{Good_C}(t) \wedge (\exists x.v_mem(x, t) \wedge \mathbf{W}[\![T]\!](x)) \quad x \notin fv(T, t)$

$\mathbf{R}[\![\{v_1, \dots, v_n\}]\!] = \text{Return}(\{v_1, \dots, v_n\})$

$\mathbf{R}[\![e_1 :: e_2]\!] =$
`Bind` $x_1 \leftarrow \mathbf{R}[\![e_1]\!]$ **in** `Bind` $x_2 \leftarrow \mathbf{R}[\![e_2]\!]$ **in**
(if `is_C(x2)` **then** `Return(v_add(x1, x2)`) **else** **Error**)

$\mathbf{R}[\![\text{from } x \text{ in } e_1 \text{ let } y = e_2 \text{ accumulate } e_3]\!] =$
`Bind` $x_1 \leftarrow \mathbf{R}[\![e_1]\!]$ **in** `Bind` $x_2 \leftarrow \mathbf{R}[\![e_2]\!]$ **in**
`res_accumulate((fun` $x \ y \rightarrow \mathbf{R}[\![e_3]\!]$), x_1, x_2)

In order to give a semantics to function applications we recall that we only consider expressions that only call labeled-pure functions, and that the body of a labeled-pure function is itself a pure expression. For each labeled-pure function definition $f(x_1 : T_1, \dots, x_n : T_n) : U\{e\}$, the model of the symbol f is the total function $\underline{f} \in$

Valueⁿ → Result such that $\underline{f}(v_1, \dots, v_n)$ is the result r such that $e\{v_1/x_1\} \dots \{v_n/x_n\} \Downarrow r$. (We know that there is a unique r such that $e\{v_1/x_1\} \dots \{v_n/x_n\} \Downarrow r$ because e is pure.) Hence, the following holds by definition:

Lemma 1

If $f(x_1 : T_1, \dots, x_n : T_n) : U\{e\}$, and e is pure, and $e\{v_1/x_1\} \dots \{v_n/x_n\} \Downarrow r$ then $\models \underline{f}(v_1, \dots, v_n) = r$.

Semantics: function application

$$\mathbf{R}[\underline{f}(e_1, \dots, e_n)] = \mathbf{Bind} \ x_1 \leftarrow \mathbf{R}[e_1] \ \mathbf{in} \ \dots \ \mathbf{Bind} \ x_n \leftarrow \mathbf{R}[e_n] \ \mathbf{in} \ \underline{f}(x_1, \dots, x_n)$$

The operational semantics preserves the logical meaning of closed pure expressions.

Lemma 2

For all closed pure expressions e and e' if $e \rightarrow e'$ then $\models \mathbf{R}[e] = \mathbf{R}[e']$.

Moreover, we have a full abstraction result for this first-order language: the equalities induced by the operational and logical semantics of closed pure expressions coincide.

Theorem 1 (Full Abstraction)

For all closed pure expressions e and e' , $\models \mathbf{R}[e] = \mathbf{R}[e']$ if and only if for all r , $e \Downarrow r \Leftrightarrow e' \Downarrow r$.

The proofs of Lemma 2 and Theorem 1 are in Appendix C.1.

We calculate the semantics of some example types from Section 2.4.

Semantics of derived forms

$$\begin{aligned} \models \mathbf{R}[e_1 == e_2] &= \mathbf{Bind} \ x_1 \leftarrow \mathbf{R}[e_1] \ \mathbf{in} \ \mathbf{Bind} \ x_2 \leftarrow \mathbf{R}[e_2] \ \mathbf{in} \ \mathbf{Return}(\mathbf{v_logical}(x_1 = x_2)) \\ \models \mathbf{F}[\mathbf{Empty}](t) &\Leftrightarrow \mathbf{false} \\ \models \mathbf{F}[\underline{e}](t) &\Leftrightarrow \mathbf{R}[e] = \mathbf{Return}(t) \\ \models \neg \mathbf{W}[\underline{T}](t) \wedge \neg \mathbf{W}[\underline{U}](t) &\Rightarrow (\mathbf{F}[\underline{T} \mid \underline{U}](t) \Leftrightarrow (\mathbf{F}[\underline{T}](t) \vee \mathbf{F}[\underline{U}](t))) \\ \models \neg \mathbf{W}[\underline{T}](t) \wedge \neg \mathbf{W}[\underline{U}](t) &\Rightarrow (\mathbf{F}[\underline{T} \ \& \ \underline{U}](t) \Leftrightarrow (\mathbf{F}[\underline{T}](t) \wedge \mathbf{F}[\underline{U}](t))) \\ \models \neg \mathbf{W}[\underline{T}](t) &\Rightarrow (\mathbf{F}[\underline{!T}](t) \Leftrightarrow \neg \mathbf{F}[\underline{T}](t)) \\ \models \bigwedge_{i \in 1..n} \neg \mathbf{W}[\underline{T}_i](\mathbf{v_dot}(t, \ell_i)) &\Rightarrow (\mathbf{F}[\{\ell_i : T_i \ i \in 1..n\}](t) \Leftrightarrow \\ &\quad \mathbf{Good_E}(t) \wedge \bigwedge_{i \in 1..n} (\mathbf{v_has_field}(\ell_i, t) \wedge \mathbf{F}[\underline{T}_i](\mathbf{v_dot}(t, \ell_i)))) \\ \models \neg \mathbf{W}[\underline{U}](t) &\Rightarrow (\mathbf{F}[\underline{(s : T \ \& \ U)}](t) \Leftrightarrow \mathbf{F}[\underline{T}](t) \wedge \mathbf{let} \ s = t \ \mathbf{in} \ \mathbf{F}[\underline{U}](t)) \end{aligned}$$

4 Algorithmic purity check

Our definition of purity defined in Section 2.3 is undecidable, so in this section we introduce a tractable property, *algorithmic purity*, on which we rely instead of purity itself in the subsequent definitions of our type systems. Algorithmic purity is defined in terms of a syntactic termination condition on function applications to avoid divergence, and a restriction on accumulate expressions to avoid non-determinism. We show, Theorem 2 below, that algorithmic purity implies purity.

We call an expression e *algorithmically pure* if and only if the following three conditions hold:

1. If e is a function application $f(e_1, \dots, e_n)$ then f is labeled-pure.
2. If e is of the form **from** x **in** e_1 **let** $y = e_2$ **accumulate** e_3 then

$$\models \mathbf{R}[\mathbf{let} \ y = e_3\{x_1/x\}\{y_1/y\} \ \mathbf{in} \ e_3\{x_2/x\}] = \\ \mathbf{R}[\mathbf{let} \ y = e_3\{x_2/x\}\{y_1/y\} \ \mathbf{in} \ e_3\{x_1/x\}]$$

(where the variables x_1 , x_2 , and y_1 do not appear free in e_3).

3. All the proper subexpressions of e are algorithmically pure (including the ones inside all refinement types contained by e).

Furthermore, we require that each labeled-pure function f has an algorithmically pure body that only calls f (directly or indirectly) on structurally smaller arguments; since termination-checking is not the focus of this paper, we omit the rather technical details, which may be found elsewhere (Giesl, 1997).

Thus condition (1) enforces termination of algorithmically pure expressions: Only labeled-pure functions can be called and if these functions are recursive, then recursive calls can only be on syntactically smaller arguments. Condition (2) only allows accumulates in an algorithmically pure expression if the order in which the elements are processed is irrelevant for the final result. In general, we call a (mathematical) function $f : X \times Y \rightarrow Y$ *order-irrelevant* if $f(x_1, f(x_2, y)) = f(x_2, f(x_1, y))$ for all x_1 , x_2 , and y . Enforcing that the semantics of the body of accumulate expressions is an order-irrelevant function is a sufficient condition for the uniqueness of evaluation results. We phrase this condition in terms of the logical semantics and check it by using the SMT solver. Order-irrelevance is less restrictive than conditions found in the literature such as associativity and commutativity (Cohen, 2006; Leino & Monahan, 2009). If f is associative and commutative then f is also order-irrelevant, but the converse fails in general.² If f is order-irrelevant, its two arguments need not even have the same type. For instance, none of the derived collection expressions from Section 2.1 is either associative or commutative, and in most of the cases the accumulator has a different type from the iterator.

Accumulate expressions are often useful inside refinements. For instance, as shown in Section 2.1, the number of elements of collections can be computed using an accumulate expression: **from** x **in** e **let** $y = 0$ **accumulate** $y + 1$. Showing that this expression is algorithmically pure boils down to showing that $f(x, y) = y + 1$ is order-irrelevant. More precisely, the (in this case trivial) proof obligation discharged by the SMT solver has the following form:

$$\models \mathbf{Bind} \ y \Leftarrow \mathbf{R}[y_1 + 1] \ \mathbf{in} \ \mathbf{R}[y + 1] = \mathbf{Bind} \ y \Leftarrow \mathbf{R}[y_1 + 1] \ \mathbf{in} \ \mathbf{R}[y + 1]$$

where $\mathbf{R}[x + y] = (\mathbf{if} \ \text{In.Integer}(x) \wedge \text{In.Integer}(y) \ \mathbf{then} \ \mathbf{Return}(\text{O.Add}(x, y)) \ \mathbf{else} \ \mathbf{Error})$

² The weaker order-irrelevance condition is sufficient in our setting because in an accumulate expression non-determinism only arises from different orders in which the elements of a collection can be processed. On the other hand, the execution model is more complicated for user-defined aggregate functions in a database (Cohen, 2006), since the database management system can exploit true parallelism, to start multiple threads of computation, which later have to be merged and their results combined.

If instead we count only the number of occurrences of a particular element z using **from** x **in** e **let** $y = 0$ **accumulate** $(x == z)?y + 1 : y$, we obtain the following (more interesting) proof obligation:

$$\begin{aligned} & \models \mathbf{Bind} \ y \leftarrow (\mathbf{if} \ \mathbf{O_EQ}(x_1, z) = \mathbf{true} \ \mathbf{then} \ \mathbf{R}[\![y_1 + 1]\!] \ \mathbf{else} \ \mathbf{Return}(y_1)) \\ & \quad \mathbf{in} \ (\mathbf{if} \ \mathbf{O_EQ}(x_2, z) = \mathbf{true} \ \mathbf{then} \ \mathbf{R}[\![y + 1]\!] \ \mathbf{else} \ \mathbf{Return}(y)) = \\ & \mathbf{Bind} \ y \leftarrow (\mathbf{if} \ \mathbf{O_EQ}(x_2, z) = \mathbf{true} \ \mathbf{then} \ \mathbf{R}[\![y_1 + 1]\!] \ \mathbf{else} \ \mathbf{Return}(y_1)) \\ & \quad \mathbf{in} \ (\mathbf{if} \ \mathbf{O_EQ}(x_1, z) = \mathbf{true} \ \mathbf{then} \ \mathbf{R}[\![y + 1]\!] \ \mathbf{else} \ \mathbf{Return}(y)) \end{aligned}$$

We show that the algorithmic purity check is a sufficient condition for purity.

Theorem 2

If e is algorithmically pure, then e is pure.

Proof

The details are given in Appendix C.2. \square

The logical semantics is defined using purity to handle the case of (labeled-pure) function applications. Given the logical semantics, we obtain algorithmic purity, a sufficient condition for purity. In the remainder of the paper we rely only on algorithmic purity.

5 Declarative type system

In this section we give a non-algorithmic type assignment relation, and prove preservation and progress properties relating it to the operational semantics. In the next section, we present algorithmic rules – the basis of our type-checker – for proving type assignment.

Each judgment of the type system is with respect to a typing *environment* E , of the form $x_1 : T_1, \dots, x_n : T_n$, which assigns a type to each variable in scope. We write \emptyset for the empty environment, $\mathit{dom}(E)$ to denote the set of variables defined by a typing environment E , and $\mathbf{F}[\![E]\!]$ for the logical interpretation of E .

Environments and their logical semantics

$$\begin{array}{ll} E ::= x_1 : T_1, \dots, x_n : T_n & \text{type environments} \\ \mathit{dom}(x_1 : T_1, \dots, x_n : T_n) = \{x_1, \dots, x_n\} \\ \mathbf{F}[\![x_1 : T_1, \dots, x_n : T_n]\!] \triangleq \mathbf{F}[\![T_1]\!](x_1) \wedge \dots \wedge \mathbf{F}[\![T_n]\!](x_n) \end{array}$$

Judgments of the declarative type system

$$\begin{array}{ll} E \vdash \diamond & \text{environment } E \text{ is well-formed} \\ E \vdash T & \text{in } E, \text{ type } T \text{ is well-formed} \\ E \vdash T <: T' & \text{in } E, \text{ type } T \text{ is a subtype of } T' \\ E \vdash e : T & \text{in } E, \text{ expression } e \text{ has type } T \end{array}$$

Global assumptions

For each function definition $f(x_1 : T_1, \dots, x_n : T_n) : U\{e_f\}$ we assume that $x_1 : T_1, \dots, x_n : T_n \vdash e_f : U$.

Rules of well-formed environments and types: $E \vdash \diamond$, $E \vdash T$

(Env Empty) $\frac{}{\emptyset \vdash \diamond}$	(Env Var) $\frac{E \vdash T \quad x \notin \text{dom}(E)}{E, x : T \vdash \diamond}$	(Type Any) $\frac{E \vdash \diamond}{E \vdash \text{Any}}$	(Type Scalar) $\frac{E \vdash \diamond}{E \vdash G}$
(Type Collection) $\frac{E \vdash T}{E \vdash T^*}$	(Type Entity) $\frac{E \vdash T}{E \vdash \{\ell : T\}}$	(Type Refine) $\frac{E, x : T \vdash e : \text{Logical} \quad e \text{ alg. pure}}{E \vdash (x : T \text{ where } e)}$	

The subtype relation is defined as logical implication between the logical semantics of well-formed types.

Rule of semantic subtyping

(Subtype) $\frac{E \vdash T \quad E \vdash T' \quad x \notin \text{dom}(E) \quad \models (\mathbf{F}[\![E]\!] \wedge \mathbf{F}[\![T]\!](x)) \Rightarrow \mathbf{F}[\![T']\!](x)}{E \vdash T <: T'}$

Rules of type assignment: $E \vdash e : T$

(Exp Singular Subsum) $\frac{E \vdash e : T \quad E \vdash [e : T] <: T'}{E \vdash e : T'}$	(Exp Var) $\frac{E \vdash \diamond \quad (x : T) \in E}{E \vdash x : T}$	(Exp Const) $\frac{E \vdash \diamond}{E \vdash c : \text{Any}}$
(Exp Eq) $\frac{E \vdash e_1 : T_1 \quad E \vdash e_2 : T_2}{E \vdash e_1 == e_2 : \text{Logical}}$	(Exp Operator) $\frac{E \vdash e_1 : T_1 \cdots E \vdash e_n : T_n \oplus : T_1, \dots, T_n \rightarrow T, \oplus \neq (==)}{E \vdash \oplus(e_1, \dots, e_n) : T}$	
(Exp Cond) $\frac{E \vdash e_1 : \text{Logical} \quad E, - : \text{Ok}(e_1) \vdash e_2 : T \quad E, - : \text{Ok}(!e_1) \vdash e_3 : T}{E \vdash (e_1 ? e_2 : e_3) : T}$		
(Exp Let) $\frac{E \vdash e_1 : T \quad E, x : T \vdash e_2 : U \quad x \notin \text{fv}(U)}{E \vdash \text{let } x = e_1 \text{ in } e_2 : U}$	(Exp Test) $\frac{E \vdash e : \text{Any} \quad E \vdash T}{E \vdash e \text{ in } T : \text{Logical}}$	
(Exp Entity) $\frac{E \vdash e_i : T_i \quad \forall i \in 1..n \quad E \vdash \diamond}{E \vdash \{\ell_i \Rightarrow e_i \}_{i \in 1..n} : \{\ell_i : T_i \}_{i \in 1..n}}$	(Exp Dot) $\frac{E \vdash e : \{\ell : T\}}{E \vdash e.\ell : T}$	
(Exp Coll) $\frac{E \vdash v_i : T \quad \forall i \in 1..n \quad E \vdash \diamond}{E \vdash \{v_1, \dots, v_n\} : T^*}$	(Exp Add) $\frac{E \vdash e_1 : T \quad E \vdash e_2 : T^*}{E \vdash (e_1 :: e_2) : T^*}$	
(Exp Acc) $\frac{E \vdash e_1 : T^* \quad E \vdash e_2 : U \quad E, x : T, y : U \vdash e_3 : U \quad x, y \notin \text{fv}(U)}{E \vdash \text{from } x \text{ in } e_1 \text{ let } y = e_2 \text{ accumulate } e_3 : U}$		

(Exp App)

$$\frac{\text{given } f(x_1 : T_1, \dots, x_n : T_n) : U\{e_f\} \\ E \vdash T_1 \quad E \vdash T_2\{e_1/x_1\} \cdots E \vdash T_n\{e_1/x_1\} \cdots \{e_{n-1}/x_{n-1}\} \quad E \vdash U\{e_1/x_1\} \cdots \{e_n/x_n\} \\ E \vdash e_1 : T_1 \quad E \vdash e_2 : T_2\{e_1/x_1\} \cdots E \vdash e_n : T_n\{e_1/x_1\} \cdots \{e_{n-1}/x_{n-1}\}}{E \vdash f(e_1, \dots, e_n) : U\{e_1/x_1\} \cdots \{e_n/x_n\}}$$

For the sake of parsimony, the conclusion $E \vdash c : \text{Any}$ of the rule (Exp Const) says only that a constant c is well-typed given that E is a well-formed environment. If $c \in K(G)$, the two alternative conclusions $E \vdash c : G$ and $E \vdash c : [c : G]$ are derivable using (Exp Singular Subsum). (The algorithmic synthesis rule (Synth Const) yields the latter.)

The rule (Exp Cond) records the appropriate test expression in the environment when typing the branches. The actual value of a type $\text{Ok}(e)$ is arbitrary, the point is simply to record that condition e holds (Gordon & Jeffrey, 2002), provided e is algorithmically pure. When e is not algorithmically pure, $\text{Ok}(e)$ is equivalent to Any .

Typed singleton types and Ok types

$$[e : T] \triangleq \begin{cases} (x : T \text{ where } x == e) & (x \notin \text{fv}(e)) \text{ if } e \text{ alg. pure} \\ T & \text{otherwise} \end{cases}$$

$$\text{Ok}(e) \triangleq \begin{cases} (x : \text{Any where } e) & (x \notin \text{fv}(e)) \text{ if } e \text{ alg. pure} \\ \text{Any} & \text{otherwise} \end{cases}$$

The rule (Exp Singular Subsum) can be seen as a combination of the following conventional rules of subsumption and singleton introduction (Aspinall, 1994).

$$\begin{array}{c} \text{(Exp Subsum)} \\ \frac{E \vdash e : T \quad E \vdash T <: T'}{E \vdash e : T'} \end{array} \quad \begin{array}{c} \text{(Exp Singleton)} \\ \frac{E \vdash e : T}{E \vdash e : [e : T]} \end{array}$$

Both these rules are derivable from (Exp Singular Subsum). In fact, we can go in the other direction too so that the type assignment relation would be unchanged were we to replace (Exp Singular Subsum) with (Exp Subsum) and (Exp Singleton). Still, the given presentation is simpler to work with because (Exp Singular Subsum) is the only rule not determined by the structure of the expression being typed.

The rule (Exp Singular Subsum) depends on the relation $E \vdash [e : T] <: T'$, which we refer to as *singular subtyping*. We illustrate (Exp Singular Subsum) and singular subtyping with regard to (Exp Const). For example, to derive that $E \vdash [42 : \text{Any}] <: \text{Integer}$ note that $\models \mathbf{F}[[42 : \text{Any}]](x) \Leftrightarrow x = 42$ and hence $\models \mathbf{F}[[42 : \text{Any}]](x) \Rightarrow \text{In.Integer}(x)$.

One might wonder why we have the separate rule (Exp Eq) for equality, rather than allowing (Exp Operator) to derive $E \vdash e_1 == e_2 : \text{Logical}$, relying on the signature $== : \text{Any}, \text{Any} \rightarrow \text{Logical}$. The reason we cannot type-check $e_1 == e_2$ in this way is because to type-check each e_i at the type Any in general requires us to use (Exp Singular Subsum), along with the fact that $E \vdash [e_i : T_i] <: \text{Any}$. When e_i is alg. pure, the syntax $[e_i : T_i]$ is short for $(x : T_i \text{ where } x == e_i)$, which to be well-formed requires us to type-check an equality $x == e_i$. To break this circularity, we provide the rule (Exp Eq) explicitly.

In the rule (Exp App), the well-formedness conditions on the argument types amount to requiring that the e_i in a dependent function application $f(e_1, \dots, e_n)$ is algorithmically pure. To form, say, $f(e)$ where e is impure, we can work around this restriction by writing **let** $x = e$ **in** $f(x)$ instead.

The following soundness property relates type assignment to the logical semantics of types and expressions. Point (1) is that the logical value of a well-typed expression satisfies the interpretation of its type as a predicate. Point (2) is that evaluating a type-test for a well-formed type cannot go wrong.

Theorem 3 (Logical Soundness)

1. If e is alg. pure and $E \vdash e : T$ then:
 - (a) $\models \mathbf{F}[[E]] \Rightarrow \text{Proper}(\mathbf{R}[[e]])$
 - (b) $\models \mathbf{F}[[E]] \Rightarrow \mathbf{F}[[T]](\text{out.v}(\mathbf{R}[[e]]))$
2. If $E \vdash U$ then $\models \mathbf{F}[[E]] \Rightarrow \forall y. \neg \mathbf{W}[[U]](y)$, for $y \notin \text{fv}(U)$.

Proof

The details of this proof are given in Appendix C.3 □

We show the safety of our type system by proving the important preservation and progress theorems (Wright & Felleisen, 1994). The details for both these proofs are given in Appendix C.4.

Theorem 4 (Preservation)

If $\emptyset \vdash e : T$ and $e \rightarrow e'$ then $\emptyset \vdash e' : T$.

Theorem 5 (Progress)

If $\emptyset \vdash e : T$ and e is not a value then $\exists e'. e \rightarrow e'$.

We conclude this section by considering the typing of the `removeNulls` function from Section 2, whose definition is as follows:

```

NullableInt  $\triangleq$  Integer | [null]
removeNulls(xs : NullableInt*) : Integer*
{ from x in xs let a = ({}:Integer*) accumulate (x!=null) ? (x :: a) : a }
    
```

The power of the Dminor type system is demonstrated in the typing of the conditional expression, i.e.

$$a : \text{Integer}^*, x : \text{NullableInt} \vdash ((x \neq \text{null}) ? (x :: a) : a) : \text{Integer}^*$$

The typing derivation is as follows (where E is the typing environment $a : \text{Integer}^*, x : \text{NullableInt}$, and e_x is the expression $x \neq \text{null}$):

$$\frac{\frac{\Pi_1}{E \vdash x \neq \text{null} : \text{Logical}} \quad \frac{\Pi_2}{E, - : \text{Ok}(e_x) \vdash (x :: a) : \text{Integer}^*} \quad \text{(Exp Var)}}{E, - : \text{Ok}(e_x) \vdash a : \text{Integer}^*} \quad \text{(Exp Cond)}$$

Derivation Π_1 is trivial; derivation Π_2 is as follows:

$$\frac{\frac{\Pi_3}{E, - : \text{Ok}(e_x) \vdash x : \text{Integer}} \quad \text{(Exp Var)}}{E, - : \text{Ok}(e_x) \vdash a : \text{Integer}^*} \quad \text{(Exp Add)}$$

Derivation Π_3 is as follows:

$$\frac{\frac{}{\text{(Exp Var)}} \quad \frac{}{\text{(Subtype)}}}{E, - : \text{Ok}(e_x) \vdash x : \text{NullableInt} \quad E, - : \text{Ok}(e_x) \vdash [x : \text{NullableInt}] <: \text{Integer}}{E, - : \text{Ok}(e_x) \vdash x : \text{Integer}} \text{(Exp Singular Subsum)}$$

As one might expect, the hard work has been delegated to semantic subtyping, i.e. the verification of the following implication.

$$(\mathbf{F}[[a : \text{Integer}^*, x : \text{NullableInt}, - : \text{Ok}(e_x)]] \wedge \mathbf{F}[[x : \text{NullableInt}]](x)) \Rightarrow \mathbf{F}[[\text{Integer}]](x)$$

The key step in verifying this implication involves the following fact about the semantics of the `Ok` type introduced into the typing environment by the `(Exp Cond)` rule (where e is any algorithmically pure expression):

$$\begin{aligned} & \mathbf{F}[- : \text{Ok}(e)] \\ \Leftrightarrow & \mathbf{F}[\text{Ok}(e)](-) \\ \Leftrightarrow & \mathbf{F}[(z : \text{Any where } e)](-) \quad z \notin \text{fv}(e) \\ \Leftrightarrow & \mathbf{F}[\text{Any}](-) \wedge \text{let } z = - \text{ in } (\mathbf{R}[e] = \text{Return}(\text{true})) \\ \Leftrightarrow & (\mathbf{R}[e] = \text{Return}(\text{true})) \end{aligned}$$

In our example e is the expression $e_x = x! = \text{null}$ and by the definition of the translation $\mathbf{R}[-]$ we obtain that $x \neq \text{null}$. By expanding the assumption that $\mathbf{F}[\text{NullableInt}](x)$ this allows us to infer that $\mathbf{F}[\text{Integer}](x)$. By expanding the encoding of typed singleton types in $\mathbf{F}[[x : \text{NullableInt}]](x)$ we deduce that $x = x$, which allows us to prove the original implication.

6 Algorithmic aspects

6.1 Optimizing the logical semantics

We built our logical semantics in Section 3 independently of the type system, and then defined our type system in Section 5 in terms of the logical semantics. Now that we have our type system, we show how to optimize the logical semantics.

Our logical semantics propagates error values so as to match the stuck expressions of our operational semantics. Tracking errors is important, but observe that when we use our logical semantics during semantic subtyping, we only ever ask whether well-formed types are related. Every expression occurring in a well-formed type is itself well-typed, and so by Theorem 3 its logical semantics is a proper value, not **Error**.

This suggests that when checking subtyping we can optimize the logical semantics given the assumption that the expressions occurring within the two types are well-typed. In particular, we can apply the following lemma to transform monadic error-checking binds into ordinary lets.

Lemma 3

If e alg. pure and $E \vdash e : T$ then $\models \mathbf{F}[[E]] \Rightarrow (\mathbf{Bind} \ x \leftarrow \mathbf{R}[[e]] \ \mathbf{in} \ t) = (\mathbf{let} \ x = \mathbf{out.v}(\mathbf{R}[[e]]) \ \mathbf{in} \ t)$.

Proof

By definition of notation, $\mathbf{Bind} \ x \leftarrow \mathbf{R}[[e]] \ \mathbf{in} \ t$ is the term **(if** $\neg \mathbf{Proper}(\mathbf{R}[[e]])$ **then else Error)** $\mathbf{let} \ x = \mathbf{out.v}(\mathbf{R}[[e]]) \ \mathbf{in} \ t$. By Theorem 3, $\models \mathbf{F}[[E]] \Rightarrow \mathbf{Proper}(\mathbf{R}[[e]])$. Hence the result. \square

The following tables present the optimized definitions used in our type-checker, and the following theorem states their correctness with respect to the error-tracking semantics of Section 3.

Optimized semantics of types: $\mathbf{F}'[[T]](t)$

$\mathbf{F}'[[\mathbf{Any}]](t) = \mathbf{true}$
 $\mathbf{F}'[[\mathbf{Integer}]](t) = \mathbf{In_Integer}(t)$
 $\mathbf{F}'[[\mathbf{Text}]](t) = \mathbf{In_Text}(t)$
 $\mathbf{F}'[[\mathbf{Logical}]](t) = \mathbf{In_Logical}(t)$
 $\mathbf{F}'[[\{\ell : T\}]](t) = \mathbf{Good_E}(t) \wedge \mathbf{v_has_field}(\ell, t) \wedge \mathbf{F}'[[T]](\mathbf{v_dot}(t, \ell))$
 $\mathbf{F}'[[T^*]](t) = \mathbf{Good_C}(t) \wedge (\forall x. \mathbf{v_mem}(x, t) \Rightarrow \mathbf{F}'[[T]](x)) \quad x \notin \mathbf{fv}(T, t)$
 $\mathbf{F}'[[x : T \ \mathbf{where} \ e]](t) = \mathbf{F}'[[T]](t) \wedge \mathbf{let} \ x = t \ \mathbf{in} \ \mathbf{V}[[e]] = \mathbf{true} \quad x \notin \mathbf{fv}(T, t)$

Optimized semantics of pure typed expressions: $\mathbf{V}[[e]]$

$\mathbf{V}[[x]] = x$
 $\mathbf{V}[[c]] = c$
 $\mathbf{V}[[\oplus(e_1, \dots, e_n)]] = \mathbf{O}_{\oplus}(\mathbf{V}[[e_1]], \dots, \mathbf{V}[[e_n]])$
 $\mathbf{V}[[e_1 ? e_2 : e_3]] = (\mathbf{if} \ \mathbf{V}[[e_1]] = \mathbf{true} \ \mathbf{then} \ \mathbf{V}[[e_2]] \ \mathbf{else} \ \mathbf{V}[[e_3]])$
 $\mathbf{V}[[\mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2]] = \mathbf{let} \ x = \mathbf{V}[[e_1]] \ \mathbf{in} \ \mathbf{V}[[e_2]]$
 $\mathbf{V}[[e \ \mathbf{in} \ T]] = \mathbf{v_logical}(\mathbf{F}'[[T]](\mathbf{V}[[e]]))$
 $\mathbf{V}[[e : T]] = \mathbf{V}[[e]]$
 $\mathbf{V}[[\{\ell_i \Rightarrow e_i \ i \in 1..n\}]] = \{\ell_i \Rightarrow \mathbf{V}[[e_i]] \ i \in 1..n\}$
 $\mathbf{V}[[e.\ell]] = \mathbf{v_dot}(\mathbf{V}[[e]], \ell)$
 $\mathbf{V}[[\{v_1, \dots, v_n\}]] = \{v_1, \dots, v_n\}$
 $\mathbf{V}[[e_1 :: e_2]] = \mathbf{v_add}(\mathbf{V}[[e_1]], \mathbf{V}[[e_2]])$
 $\mathbf{V}[[\mathbf{from} \ x \ \mathbf{in} \ e_1 \ \mathbf{let} \ y = e_2 \ \mathbf{accumulate} \ e_3]] =$
 $\quad \mathbf{v_accumulate}((\mathbf{fun} \ x \ y \rightarrow \mathbf{V}[[e_3]]), \mathbf{V}[[e_1]], \mathbf{V}[[e_2]])$

Optimized semantics of environments

$\mathbf{F}'[[x_1 : T_1, \dots, x_n : T_n]] \triangleq \mathbf{F}'[[T_1]](x_1) \wedge \dots \wedge \mathbf{F}'[[T_n]](x_n)$

Theorem 6 (Soundness and Completeness of Optimized Semantics)

1. If $E \vdash \diamond$ then $\models \mathbf{F}[[E]] \Leftrightarrow \mathbf{F}'[[E]]$.
2. If $E \vdash T$ and $x \notin \mathbf{dom}(E)$ then: $\models \mathbf{F}'[[E]] \Rightarrow (\mathbf{F}[[T]](x) \Leftrightarrow \mathbf{F}'[[T]](x))$.
3. If $E \vdash e : T$ then: $\models \mathbf{F}'[[E]] \Rightarrow (\mathbf{R}[[e]] = \mathbf{Return}(\mathbf{V}[[e]]))$.

Proof

The proof is by simultaneous induction on the derivations of $E \vdash \diamond$ and $E \vdash T$ and $E \vdash e : T$, with appeal to Theorem 3 and Lemma 3. \square

6.2 Bidirectional typing rules

The Dminor type-checker is implemented as a *bidirectional* type system (Pierce & Turner, 2000). The key concept of bidirectional type systems is that there are two typing relations, one for type *checking* and another for type *synthesis*. The chief characteristic of these relations is that they are local in the sense that type information is passed between adjacent nodes in the syntax tree. This is an important feature, not least because it makes type error reporting easy – a disadvantage of languages that use ML-style inference (Lerner et al., 2007). Moreover, bidirectional type systems are simple to implement, predictable for programmers, and expressive; for example, the type system for C# can be defined as a bidirectional type system (Bierman et al., 2007), and several dependently typed languages have bidirectional type systems (Knowles et al., 2007; Lovas & Pfenning, 2007).

Judgments of the algorithmic type system

$E \vdash e \rightarrow T$	in E , expression e synthesizes type T
$E \vdash e \leftarrow T$	in E , expression e checks against type T
$E \triangleright \diamond$	environment E is alg. well-formed
$E \triangleright T$	in E , type T is alg. well-formed
$E \triangleright S <: T$	in E , type S is alg. a subtype of type T

The reader will recall that the rules characterizing well-formed environments and types in Section 5 made use of the *declarative* typing relation in rule (Type Refine). We thus need to define algorithmic versions of these rules that make use of the *bidirectional* type system.

Rules of algorithmic well-formedness: $E \triangleright \diamond$, $E \triangleright T$

(Alg. Env Empty)	(Alg. Env Var)	
$\frac{}{\emptyset \triangleright \diamond}$	$\frac{E \triangleright T \quad x \notin \text{dom}(E)}{E, x : T \triangleright \diamond}$	
(Alg. Type Any)	(Alg. Type Scalar)	(Alg. Type Collection)
$\frac{E \triangleright \diamond}{E \triangleright \text{Any}}$	$\frac{E \triangleright \diamond}{E \triangleright G}$	$\frac{E \triangleright T}{E \triangleright T^*}$
(Alg. Type Entity)	(Alg. Type Refine)	
$\frac{E \triangleright T}{E \triangleright \{\ell : T\}}$	$\frac{E, x : T \vdash e \leftarrow \text{Logical} \quad e \text{ alg. pure}}{E \triangleright (x : T \text{ where } e)}$	

We also make use of the optimized semantics from Section 6.1 to define the algorithmic semantic subtyping rule.

Rule of algorithmic semantic subtyping

(Alg. Subtype)

$$\frac{E \triangleright T \quad x \notin \text{dom}(E) \quad \models (\mathbf{F}'\llbracket E \rrbracket \wedge \mathbf{F}'\llbracket T \rrbracket(x)) \Rightarrow \mathbf{F}'\llbracket T' \rrbracket(x)}{E \triangleright T <: T'}$$

Rules of type synthesis: $E \vdash e \rightarrow T$

(Synth Var)

$$\frac{E \triangleright \diamond \quad (x : T) \in E}{E \vdash x \rightarrow [x : T]}$$

(Synth Const)

$$\frac{E \triangleright \diamond}{E \vdash c \rightarrow [c : \text{typeof}(c)]}$$

(Synth Operator)

$$\frac{E \vdash e_i \leftarrow T_i \quad \forall i \in 1..n \quad \oplus : T_1, \dots, T_n \rightarrow T}{E \vdash \oplus(e_1, \dots, e_n) \rightarrow [\oplus(e_1, \dots, e_n) : T]}$$

(Synth Cond)

$$\frac{E \vdash e_1 \leftarrow \text{Logical} \quad E, - : \text{Ok}(e_1) \vdash e_2 \rightarrow T_2 \quad E, - : \text{Ok}(!e_1) \vdash e_3 \rightarrow T_3}{E \vdash (e_1 ? e_2 : e_3) \rightarrow (\text{if } e_1 \text{ then } T_2 \text{ else } T_3)}$$

(Synth Let)

$$\frac{E \vdash e_1 \rightarrow T_1 \quad E, x : T_1 \vdash e_2 \rightarrow T_2 \quad E \vdash T_2\{e_1/x\}}{E \vdash \text{let } x = e_1 \text{ in } e_2 \rightarrow T_2\{e_1/x\}}$$

(Synth Test)

$$\frac{E \vdash e \leftarrow \text{Any} \quad E \triangleright T}{E \vdash e \text{ in } T \rightarrow [e \text{ in } T : \text{Logical}]}$$

(Synth Ascribe)

$$\frac{E \vdash e \leftarrow T}{E \vdash (e^T) \rightarrow T}$$

(Synth Entity)

$$\frac{E \vdash e_1 \rightarrow T_1 \cdots E \vdash e_n \rightarrow T_n \quad E \triangleright \diamond}{E \vdash \{\ell_i \Rightarrow e_i \quad i \in 1..n\} \rightarrow \{\ell_1 : T_1\} \& \cdots \& \{\ell_n : T_n\}}$$

(Synth Dot)

$$\frac{E \vdash e \rightarrow T \quad \text{norm}(T) = D \quad D.\ell \rightsquigarrow U}{E \vdash e.\ell \rightarrow [e.\ell : U]}$$

(Synth Coll)

$$\frac{E \vdash v_i \rightarrow T_i \quad \forall i \in 1..n \quad E \triangleright \diamond}{E \vdash \{v_1, \dots, v_n\} \rightarrow (T_1 \mid \dots \mid T_n)^*}$$

(Synth Add)

$$\frac{E \vdash e_1 \rightarrow T_1 \quad E \vdash e_2 \rightarrow T_2 \quad \text{norm}(T_2) = D_2 \quad D_2.\text{Items} \rightsquigarrow U_2}{E \vdash e_1 :: e_2 \rightarrow ([e_1 : T_1] \mid U_2)^*}$$

(Synth Acc)

$$\frac{E \vdash e_1 \rightarrow T_1 \quad \text{norm}(T_1) = D_1 \quad D_1.\text{Items} \rightsquigarrow U_1 \quad E \vdash e_2 \rightarrow T_2 \quad E, x : U_1, y : T_2 \vdash e_3 \leftarrow T_2}{E \vdash \text{from } x \text{ in } e_1 \text{ let } y = e_2 \text{ accumulate } e_3 \rightarrow T_2}$$

(Synth App)

$$\frac{\text{given } f(x_1 : T_1, \dots, x_n : T_n) : U\{e_f\} \quad E \triangleright T_1 E \triangleright T_2\{e_1/x_1\} \cdots E \triangleright T_n\{e_n/x_n\} \quad E \triangleright U\{e_1/x_1\} \cdots \{e_n/x_n\} \quad E \vdash e_1 \leftarrow T_1 \quad E \vdash e_2 \leftarrow T_2\{e_1/x_1\} \cdots E \vdash e_n \leftarrow T_n\{e_1/x_1\} \cdots \{e_{n-1}/x_{n-1}\}}{E \vdash f(e_1, \dots, e_n) \rightarrow U\{e_1/x_1\} \cdots \{e_n/x_n\}}$$

The rules (Synth Var) and (Synth Const) follow the work of Aspinall, (1994) and yield singleton types for all variables and constants, where the function *typeof* returns the type of a given constant; let $typeof(c) = G$ if and only if $c \in K(G)$. Rule (Synth Entity) uses intersection types to encode multiple-field entity types (see Section 2.4).

The (Synth Cond) rule synthesizes a type **if** e_1 **then** T_2 **else** T_3 , defined below, for the conditional expression $e_1 ? e_2 : e_3$. The synthesized type is the union of the two types synthesized for the branches where we additionally record the test expression in the type (if it is algorithmically pure), which allows for more precise typing.

Encoding of conditional types

$$\text{if } e \text{ then } T \text{ else } U \triangleq \begin{cases} (- : T \text{ where } e) \mid (- : U \text{ where } !e) & \text{if } e \text{ alg. pure} \\ T \mid U & \text{otherwise} \end{cases}$$

The rule (Synth Let) faces the problem that the bound variable x should not escape into the result type T_2 , and does so by substituting e_1 for x in T_2 . In case $x \in fv(T_2)$ and e_1 is not pure, the rule does not apply, as the result type is not well-defined. In this case, the programmer needs to insert a type-ascription to remove the bound variable explicitly. Similarly, the rule (Synth App) for an application $f(e_1, \dots, e_n)$ returns a type possibly containing the expressions e_1, \dots, e_n . The rule is not applicable if these expressions are impure and occur in the result type. The programmer can work around this limitation by using a let-expression to compute the value of any impure expression before making the call to f . In practice we found that conversion to A-normal form can further improve the precision of our algorithmic type system, since purity checking is then done at a finer granularity and thus more singleton types are synthesized.

The (Synth Ascribe) rule allows the user to provide hints to the type-checker in the form of type annotations (e^T). Such type annotations are not part of the core language and have no operational significance, and are necessary in case the type-checker cannot infer the loop invariants of accumulate expressions. Although in the current presentation monadic bind expressions (Buneman *et al.*, 1995) are encoded using **accumulate** (see Section 2.1), the Dminor type-checker infers loop invariants for **bind** (and also for LINQ queries, Meijer *et al.*, 2007) using an additional (Synth Bind) rule, which exploits the encoding of collection quantifiers as refinement types (see Section 2.4).

Inferring loop invariants for bind

(Synth Bind)

$$E \vdash e_1 \rightarrow T_1 \quad \text{norm}(T_1) = D_1 \quad D_1.\text{Items} \rightsquigarrow U_1$$

$$E, x : U_1 \vdash e_2 \rightarrow T_2 \quad \text{norm}(T_2) = D_2 \quad D_2.\text{Items} \rightsquigarrow U_2$$

$$\text{if } E \triangleright U_2 \text{ then } T = U_2, \text{ otherwise } T = (y : \text{Any where } (e_1 \text{ in exists}(x : U_1)(y \text{ in } U_2)))$$

$$E \vdash (\text{bind } x \leftarrow e_1 \text{ in } e_2) \rightarrow T^*$$

As explained above, type annotations are also sometimes necessary for let expressions and function applications.

In several of the type synthesis rules we need to inspect components of intermediate types. In simple type systems this is straightforward as one can rely on the syntactic structure of types, but for rich type systems, such as the one of Dminor, this is not possible. In other dependently typed languages, either the programmer is required to insert casts to force the type into the appropriate syntactic shape (Xi & Pfenning, 1999), or types are first executed until a normal form is reached (Aspinall & Hofmann, 2005). Unfortunately, neither approach is acceptable in Dminor: the former forces too many casts on the programmer, and the latter is not feasible because refinements often refer to rather large data sets. One pragmatic possibility is to attempt type normalization but place some ad hoc bound on evaluation (Knowles *et al.*, 2007). As an alternative, we define a disjunctive normal form (DNF) for types, along with a normalization function, *norm*, for translating types into DNF, and procedures for extracting type information from DNF types. In practice, this approach works well.³

Disjunctive normal form types (DNF) and Normalization

$D ::= R_1 \mid \dots \mid R_n$	normal disjunction (Empty if $n = 0$)
$R ::= x : C \text{ where } e$	normal refined conjunction
$C ::= A_1 \& \dots \& A_n$	normal conjunction (Any if $n = 0$)
$A ::= G \mid T^* \mid \{\ell : T\}$	atomic type

$norm(\text{Any}) \triangleq x : \text{Any} \text{ where true}$

$norm(G) \triangleq x : G \text{ where true}$

$norm(T^*) \triangleq x : T^* \text{ where true}$

$norm(\{\ell : T\}) \triangleq x : \{\ell : T\} \text{ where true}$

$norm(x : T \text{ where } e) \triangleq \prod_{i=1}^n Conj_{DD}(x_i : C_i \text{ where } e_i, norm_r(x : C_i \text{ where } e))$
 where $\prod_{i=1}^n (x_i : C_i \text{ where } e_i) = norm(T)$

$norm_r(x : C \text{ where } x \text{ in } T) \triangleq norm(C \& T) \quad \text{where } x \notin fv(T)$

$norm_r(x : C \text{ where } e_1 \parallel e_2) \triangleq norm_r(x : C \text{ where } e_1) \mid norm_r(x : C \text{ where } e_2)$

$norm_r(x : C \text{ where } e_1 \&\& e_2) \triangleq Conj_{DD}(norm_r(x : C \text{ where } e_1), norm_r(x : C \text{ where } e_2))$

$norm_r(x : C \text{ where } e) \triangleq (x : C \text{ where } e) \quad \text{otherwise}$

$Conj_{DD}((R_1 \mid \dots \mid R_n), D) \triangleq Conj_{RD}(R_1, D) \mid \dots \mid Conj_{RD}(R_n, D)$

$Conj_{RD}(R, (R_1 \mid \dots \mid R_n)) \triangleq Conj_{RR}(R, R_1) \mid \dots \mid Conj_{RR}(R, R_n)$

$Conj_{RR}(x_1 : C_1 \text{ where } e_1, x_2 : C_2 \text{ where } e_3) \triangleq y : C_1 \& C_2 \text{ where } e_1\{y/x_1\} \&\& e_2\{y/x_2\}$
 where $y \notin fv(C_1, C_2, e_1, e_2)$

Normalization is defined using two functions: *norm*, which normalizes a type, and *norm_r*, which normalizes a refinement type based on the structure of the refinement expression. We make use of helper functions to build DNF types, principally the function, *Conj_{DD}*, which returns in DNF the conjunction of two disjunction types.

³ A further alternative would be to embed the normalization process into subtyping (Castagna & Chen, 2001). We leave this for future work.

We define *partial* functions to extract field and item types from normalized entity and collection types, respectively.

Extraction of field type: $D.\ell \rightsquigarrow U$

(Field Disj) $\frac{R_i.\ell \rightsquigarrow U_i \quad \forall i \in 1..n}{(R_1 \mid \dots \mid R_n).\ell \rightsquigarrow (U_1 \mid \dots \mid U_n)}$	(Field Refine) $\frac{C.\ell \rightsquigarrow U}{(x : C \text{ where } e).\ell \rightsquigarrow U}$
(Field Conj) $\frac{S = \{U_i \mid A_i.\ell \rightsquigarrow U_i\} \neq \emptyset}{(A_1 \ \& \ \dots \ \& \ A_n).\ell \rightsquigarrow (\& S)}$	(Field Atom) $\frac{}{\{\ell : T\}.\ell \rightsquigarrow T}$

The (Field Disj) rule requires that for every disjunct R_i there is a U_i such that $R_i.\ell \rightsquigarrow U_i$. In contrast, the (Field Conj) rule requires only that there is at least one conjunct A_i for which there is a U_i such that $A_i.\ell \rightsquigarrow U_i$.

Extraction of item type: $D.\text{Items} \rightsquigarrow U$

(Items Disj) $\frac{R_i.\text{Items} \rightsquigarrow U_i \quad \forall i \in 1..n}{(R_1 \mid \dots \mid R_n).\text{Items} \rightsquigarrow (U_1 \mid \dots \mid U_n)}$	(Items Refine) $\frac{C.\text{Items} \rightsquigarrow U}{(x : C \text{ where } e).\text{Items} \rightsquigarrow U}$
(Items Conj) $\frac{S = \{U_i \mid A_i.\text{Items} \rightsquigarrow U_i\} \neq \emptyset}{(A_1 \ \& \ \dots \ \& \ A_n).\text{Items} \rightsquigarrow (\& S)}$	(Items Atom) $\frac{}{(T*).\text{Items} \rightsquigarrow T}$

Rules of type checking: $E \vdash e \leftarrow T$

(Swap) $\frac{E \vdash e \rightarrow T \quad E \triangleright [e : T] < : T'}{E \vdash e \leftarrow T'}$	
(Check Cond) $\frac{E \vdash e_1 \leftarrow \text{Logical} \quad E, - : \text{Ok}(e_1) \vdash e_2 \leftarrow T \quad E, - : \text{Ok}(!e_1) \vdash e_3 \leftarrow T}{E \vdash e_1 ? e_2 : e_3 \leftarrow T}$	
(Check Let) $\frac{E \vdash e_1 \rightarrow T \quad E, x : T \vdash e_2 \leftarrow U \quad x \notin \text{fv}(U)}{E \vdash \text{let } x = e_1 \text{ in } e_2 \leftarrow U}$	(Check Dot) $\frac{E \vdash e \leftarrow \{\ell : T\}}{E \vdash e.\ell \leftarrow T}$

The (Swap) rule tests for singular subsumption. In our implementation we apply this rule only if the expression to be type-checked is not a conditional, let-expression, or a field selection. Typically (for example, SAGE, Knowles *et al.*, 2007), the type-checking relation for a bidirectional type system consists of a single rule of the form:

$$\frac{E \vdash e \rightarrow S \quad E \triangleright S < : T}{E \vdash e \leftarrow T}$$

However, we have found in practice that in the cases where the expression is a conditional or a let-expression, we get better precision of type-checking by following Pierce & Turner, (2000) and passing the type through to the subexpressions, as shown in the (Check Cond) and (Check Let) rules. Similarly, we can pass through an entity type in the (Check Dot) rule.

Theorem 7 (Soundness of Algorithmic Type System)

1. If $E \triangleright \diamond$ then $E \vdash \diamond$.
2. If $E \triangleright T$ then $E \vdash T$.
3. If $E \triangleright S <: T$ and $E \vdash S$ then $E \vdash S <: T$.
4. If $E \vdash e \rightarrow T$ then $E \vdash e : T$.
5. If $E \vdash e \leftarrow T$ then $E \vdash e : T$.

Proof

The details of this proof are in Appendix C.5. \square

7 Exploiting SMT models

SMT solvers, such as Z3, can produce a potential model in case they fail to prove the validity of a proof obligation (that is, when they show the satisfiability of its negation, or when they give up). In our case such models can be automatically converted into assignments mapping program variables to Dminor values. Because of the inherent incompleteness of the SMT solver⁴ and of the axiomatization we feed to it, the obtained assignment is not guaranteed to be correct. However, given a way to validate assignments, one can use the correct ones to provide very precise counterexamples when type-checking fails, and to find inhabitants of types statically or dynamically in a way that amounts to a new style of constraint logic programming (Jaffar & Maher, 1994).

7.1 Precise counterexamples to type-checking

The type-checking algorithm from Section 6.2 crucially relies on subtyping, as in the rule (Swap), and our algorithmic semantic subtyping relation $E \triangleright T <: T'$ produces proof obligations of the form

$$\models (\mathbf{F}'[E] \wedge \mathbf{F}'[T](x)) \Rightarrow \mathbf{F}'[T'](x)$$

for some fresh variable x . If the SMT solver fails to prove such an obligation, it produces a potential model, from which we can extract an assignment σ mapping x and all variables in E to Dminor values. To verify that σ is a valid counterexample, we check the following three conditions:

1. $E \triangleright T$ and $E \triangleright T'$;

⁴ Other than background theories with a non-recursively enumerable set of logical consequences such as integer arithmetic, other sources of incompleteness in SMT solvers are quantifiers (which are usually heuristically instantiated) and user-defined time-outs. All these three sources of incompleteness affect Dminor.

2. $(y\sigma \text{ in } U\sigma) \rightarrow^* \text{true}$, for all $(y : U) \in E$;
3. $(x\sigma \text{ in } (T \ \&!T')\sigma) \rightarrow^* \text{true}$.

Condition (1) enforces that we only evaluate pure expressions, therefore ensuring termination and confluence of the reduction. Condition (2) enforces that the values for all variables in E have their corresponding (possibly dependent) types. Condition (3) checks whether the value assigned to x by σ is an element of T but not an element of T' . If these three checks succeed, σ is a valid counterexample to typing that we display to the user.

Since the type-checker is itself over-approximating, there is no guarantee that an expression e that fails to type-check is going to get stuck when evaluated. The best we might do is to evaluate $e\sigma$ for a fixed number of steps, a fixed number of times (remember that e can be non-deterministic), searching for a counterexample trace we can additionally display to the user.

7.2 Finding elements of types statically

Type emptiness can be phrased in terms of subtyping as $E \vdash T <: \text{Empty}$, or equivalently $\models \neg(\mathbf{F}[\![E]\!] \wedge \mathbf{F}[\![T]\!](x))$ for some fresh x . We additionally check that $\mathbf{F}[\![E]\!]$ is satisfiable (and the model the SMT solver produces is a correct one) to exclude the case that the environment is inconsistent and therefore any subtyping judgment holds vacuously. Hence, we can detect empty types during type-checking and issue a warning to the user if an empty type is found. This is useful, since one can make mistakes when writing types containing complicated constraints. Moreover, if the SMT solver cannot prove that a type is empty, we again obtain an assignment σ , which we can validate as in Section 7.1. If validation succeeds, we know that $x\sigma$ is an element of $T\sigma$, and we can display this information if the user hovers over a type.

7.3 Finding elements of types dynamically

We can use the same technique to find elements of types dynamically. We augment the calculus with a new primitive expression **elementof** T (not present in the M language) which tries to find an inhabitant of T . If successful the expression returns such a value, otherwise it returns **null**. (We can always choose T so that **null** is not a member so that returning **null** unambiguously signals that no member of T was found.)

Operational semantics for finding elements of types

elementof $T \rightarrow v$ where $v \text{ in } T \rightarrow^* \text{true}$
elementof $T \rightarrow \text{null}$

Finding elements of types is actually simpler to do dynamically than statically: at run-time all variables inside types have already been substituted by values, so there are fewer checks to perform.

The outcome of **elementof** T is in general non-deterministic, and depends in practice on the computational power and load of the system as well as on the timeout used when calling the SMT solver. Because of this **elementof** T , expressions are considered algorithmically impure, and therefore cannot appear inside types.

Typing rules for elementof

$\frac{}{(Exp\ elementof)} \quad E \vdash T$	$\frac{}{(Synth\ elementof)} \quad E \vdash T$
$E \vdash \mathbf{elementof}\ T : (T \mid [\mathbf{null}])$	$E \vdash \mathbf{elementof}\ T \rightarrow (T \mid [\mathbf{null}])$

The new **elementof** T construct enables a form of constraint programming in Dminor, in which we iteratively change the constraints inside types in order to explore a large state space. For instance, the following recursive function computes all correct configurations of a complex system when called with the empty collection as argument. Correctness is specified by some type GoodConfig.

```
allGoodConfigs(avoid : GoodConfig*) : GoodConfig* {
  let m = elementof (GoodConfig where !(value in avoid)) in
  (m == null) ? {} : (m :: (allGoodConfigs(m :: avoid)))
}
```

Programming in this purely declarative style can be appealing for rapid prototyping or other tasks where efficiency is not the main concern. One only needs to specify *what* has to be computed in the form of a type. It is up to the SMT solver to use the right decision procedures and heuristics to perform the computation. If this fails or is too slow, one can instead implement the required functionality manually. There is little productivity loss in this case since the types one has already written will serve as specification for the code that needs to be written manually.

8 Implementation

Our prototype Dminor implementation is approximately 2,700 lines of F# code, excluding the lexer and parser. Our type-checker implements the optimized logical semantics from Section 6.1, and the bidirectional typing rules from Section 6.2. We use Z3 (de Moura & Bjørner, 2008) to discharge the proof obligations generated by semantic subtyping. Together with the proof obligations we feed to Z3 a 500 line axiomatization of our intended model in SMT-LIB format (Ranise & Tinelli, 2006), which uses the theories of integers, datatypes, and extensional arrays (see Appendix B). Our Coq formalization is just over 4,000 lines of Coq, out of which the definition of the intended model of Dminor and the proof of its well-definedness are about 2,000 lines (see Appendix A). The rest encompasses formalizations of our definitions and mechanized versions of some of the proofs in the paper (Lemmas 12–14, 19, 22, 23, 43, 44, and Theorem 5).

We have tested our type-checker on a test suite consisting of about 130 files, some type-correct and some type-incorrect, some hand-crafted by us and some transliterated from the M preliminary release. Even without serious optimization the type-checker is fast. Checking each of the 130 files in our test suite on a typical

laptop takes from under 1 second (for just startup and parsing) to around 3 seconds (for type-checking a 175 lines long interpreter for while-programs – see Section 1.1 – that discharges more than 300 proof obligations). Also, our experience with Z3 has been very positive so far – while it is possible to craft subtyping tests that cannot be efficiently checked,⁵ Z3 has performed very well on the idioms in our test suite. Still, we cannot draw firm conclusions until we have studied bigger examples.

We have also implemented the techniques for exploiting SMT solver models described in Section 7. We built a plugin for the Microsoft Intellipad text editor (Microsoft Corporation 2009) that displays precise counterexamples to typing, flags empty types, and otherwise displays one element of each type defined in the code. Moreover, our interpreter for Dminor supports `elementof` for dynamically generating instances of types (Section 7.3). This works well for simple constraints involving equalities, datatypes and simple arithmetic, and types that are not too deeply nested. However, scaling this up to arbitrary Dminor types is a challenge that will require additional work, as well as further progress in SMT solvers.

9 Related work

While Dminor’s combination of refinement types and type-tests is new and highly expressive, it builds upon a large body of related work on advanced type systems. Refinement types have their origins in early work in theorem proving systems and specification languages, such as subset types in constructive type theory (Nordström & Petersson, 1983), set comprehensions in VDM (Jones, 1986), and predicate subtypes in PVS (Rushby *et al.*, 1998). In PVS, constraints found when checking predicate subtypes become proof obligations to be proved interactively. More recently, Sozeau, (2006) extends Coq with subset types; as in PVS the proofs of subset type membership have to be constructed using tactics.

Pratt, (1983) argued for a semantic notion of “predicate types,” where objects intrinsically belong to many types. His proposed language, Viron, has an early notion of refinement type. Freeman & Pfenning, (1991) extended ML with a form of refinement type, and Xi & Pfenning, (1999) considered applications of dependent types in an extension of ML. In both of these systems, decidability of type checking is maintained by restricting which expressions can appear in types. Lovas & Pfenning, (2007) presented a bidirectional refinement type system for LF, where a restriction on expressions leads to an expressive yet decidable type system.

Other work has combined refinement types with syntactic subtyping (Bengtson *et al.*, 2008; Rondon *et al.*, 2008) but none includes type-test in the refinement language. Closest to our type system is the work of Flanagan (2006) on hybrid types and SAGE (Knowles *et al.*, 2007). SAGE also uses an SMT solver to check the validity of refinements but not for subtyping (checked by traditional syntactic techniques), and does not allow type-test expressions in refinements. However, SAGE supports a dynamic type and employs a particular form of hybrid

⁵ Z3 gets at most 1 second for each proof obligation by default.

type-checking (Flanagan, 2006; Knowles & Flanagan, 2010) that allows particular expressions to have their type-check deferred until run-time. The idea of hybrid types is to strike a balance between runtime checking of contracts as in Eiffel (Meyer, 1992) and Racket (Findler & Felleisen, 2002), and static typing. Compared with purely static typing, this can reduce the number of false alarms generated by type-checking.

In spite of early work on semantic subtyping by Aiken & Wimmers, (1993) and Damm, (1994), most programming and query languages instead use a *syntactic* notion of subtyping. This syntactic approach is typically formalized by an inductively or co-inductively defined set of rules (Pierce, 2002). Unfortunately, deriving an algorithm from such a set of rules can be difficult, especially for advanced features such as intersection and union types (Dunfield & Pfenning, 2004; Dunfield, 2007).

Although by definition semi-structured data (such as from the SSD model (Abiteboul *et al.*, 2000) or, more recently, XML and JSON) has no schema, Buneman & Pierce, (1999) show that it can be checked using a flexible enough type system. They propose a combination of collection types, record types, and untagged union types, along with a sophisticated notion of structural subtyping, including certain distributivity axioms.

The introduction of XML and XML query languages led to renewed (practical) interest in semantic subtyping. In the context of XML documents, there is a natural generalization of document type definitions where the structures in XML documents can be described using regular expression types. These types capture and generalize regular expression notation (such as *, ?, and |) and subtyping becomes inclusion between the sets of documents denoted by two regular expression types.⁶ Hosoya and Pierce (2003) first defined such a type system for XML (Hosoya *et al.*, 2000) and an XML-processing language, XDuce, based on this type system. Frisch *et al.*, (2008) extended semantic subtyping to function types and propositional types, with type-test, resulting in the language CDuce (Benzaken *et al.*, 2003). (An excellent overview of the use of semantic subtyping in the context of querying XML documents was given by Castagna, 2005.) In the end, the XQuery working group resorted to a more conventional pure named type system (Siméon & Wadler, 2003) with a simpler notion of subtyping based on ordinary regular expression inclusion (as opposed to XDuce's use of tree regular expressions). Neither XDuce nor CDuce provides general refinement types, and their subtype algorithm is purpose-built. Genevès *et al.*, (2007) consider a related problem of XML path containment. They translate XPath expressions and XML regular tree types into a particular logic and hence containment becomes implication. They use binary decision diagrams to check satisfiability; however, their type system does not provide general refinement types.

CDuce allows expressions to be pattern-matched against types and statically detects if a pattern-matching expression is non-exhaustive or a branch is unreachable. If this is the case, a counterexample XML document is generated that exhibits the problem. CDuce can also generate a counterexample document when subtyping fails, and issues warnings if empty types are detected. These tasks are much simpler

⁶ More precisely, regular expression types correspond exactly to tree automata and thus subtyping reduces to the inclusion problem between tree automata.

in CDuce than they are in our setting, since we additionally have to deal with general refinement types. In particular, the models produced by the SMT solver are not guaranteed to be real counterexamples, so we perform additional validation as explained in Section 7.

Complete functional synthesis (Kuncak *et al.*, 2010) is a recent technique closely related to our **elementof** expressions. It involves generating specialized decision procedures to find concrete values that satisfy predicates expressed as Boolean expressions in a predictable way, while our **elementof** expressions find concrete values that satisfy predicates expressed as refinement types by calling an SMT solver. It would be interesting to investigate whether we can compile **elementof** expressions into code implementing specialized decision procedures following the techniques of Kuncak *et al.* (2010).

X10 (Saraswat *et al.*, 2008) is an object-oriented language that supports refinement types. A class C can be refined with a constraint c on the immutable state of C , resulting in a type written $C(:c)$. The base language supports only simple equality constraints but further constraints can be added and multiple constraint solvers can be integrated into a compiler. In comparison with Dminor, X10 uses a mixture of semantic and syntactic subtyping, while its constraint language (Saraswat *et al.*, 2008, Section 2.11) does not support type-test expressions.

Soft typing systems (Cartwright & Fagan, 1991; Aiken *et al.*, 1994; Wright & Cartwright, 1997) infer types that represent program invariants, including shapes of S-expressions, but are not value-dependent. Typed Scheme (Tobin-Hochstadt & Felleisen, 2008, 2010) makes use of shallow type-test expressions, union types and notions of visible and latent predicates to type-check Scheme programs. Typed Scheme records information from previous conditional expressions in a similar way as we do in our rule (Exp Cond). It would be interesting to see if these idioms can be internalized in the Dminor type system using refinements – preliminary experiments are encouraging: Dminor can already handle all the first-order challenge examples introduced in the paper of Tobin-Hochstadt & Felleisen, (2010). We list our type-checked translation of some of these examples below.

```
f(x:Integer | Text) : Integer {
  (x in Integer) ? (x + 1) : string_length(x) }
```

```
g(x : Integer | Text | Logical) : Any {
  (let temp = (x in Integer) in
    ((temp) ? temp :(x in Text)))
  ?f(x)
  :0 }
```

```
h(input : (Integer | Text), extra : {fst: Any; snd: Any;}) : Integer {
  ((input in Integer) && (extra.fst in Integer)) ?
    (input + extra.fst)
  : ((extra.fst in Integer) ?
    (string_length(input) + extra.fst)
    : 0) }
```

More recently, Greenberg *et al.*, (2010) have considered the use of refinement types combined with dynamic checks and first-class blame as a higher order contracts language in the sense of Findler & Felleisen, (2002). It remains for future work to similarly consider an extension of Dminor with first-class blame as a contracts language. Some other prior work on dependent type systems has specifically targeted correct access to union types in COBOL (Komondoor *et al.*, 2005) and in C (Jhala *et al.*, 2007).

PADS (Fisher *et al.*, 2006) develops a type theory for ad hoc data formats, such as system traces, together with a rich range of tools for learning such formats and integrating into existing programming languages. The PADS type theory has refinement types, dependent pairs, and intersection types, but no type-test. There is a syntactic notion of type equivalence, but not subtyping. Dminor would be a useful language for programming transformations on data parsed using PADS, as our type system would enforce the constraints in PADS specifications, and hence guarantee statically that transformed data remain well-formed. Existing interfaces of PADS to C or to OCaml do not offer this guarantee.

DVerify (Backes *et al.*, 2011) is a recent tool that verifies Dminor programs by translating them into a standard while language and then using Boogie (Barnett *et al.*, 2005) for generating verification conditions. DVerify directly uses our logical semantics from Section 3 to generate assertions in the while-program that faithfully represent the typing constraints in the original Dminor program. Experimental evidence suggests that DVerify achieves very similar precision and efficiency compared to our prototype type-checker.

We do not consider type inference for Dminor; we assume that all function definitions have explicit type signatures. There has been considerable recent progress in inference algorithms for refinement types (Rondon *et al.*, 2008; Kawaguchi *et al.*, 2009; Unno & Kobayashi, 2009; Terauchi, 2010; Jhala *et al.*, 2011), some of which may be applicable to inferring type signatures for Dminor functions.

SMT solvers are widely used to find concrete inputs to imperative programs by using forms of symbolic execution (King, 1976). Our use of an SMT solver to find concrete values of types is in the same spirit. Both ideas amount to asking the SMT solver to find concrete counterexamples to formulas. To relate these approaches, it would be interesting for future work to consider a symbolic execution of our operational semantics for type-test, and to compare the resulting formulas with our direct interpretation of refinement types as formulas.

10 Conclusions

We have described Dminor, a simple, first-order functional language for data processing that features an especially expressive type system. The novel combination of refinement types and type-test allows us to encode a rich variety of typing idioms; for example, intersection, union, negation, singleton, nullable, variant, and algebraic types are all derivable.

The main contribution of this paper is a technique to type-check Dminor programs *statically*: We combine the use of a bidirectional type system with the use of an

SMT solver to perform semantic subtyping. Previous type systems have either devised special purpose algorithms for semantic subtyping, or used theorem provers only for refinement types. As far as we are aware, our use of an SMT solver to determine Dminor's very general notion of semantic subtyping is novel. We have implemented our type system in $F^\#$ using the Z3 SMT solver. We consider that SMT solvers are now of sufficient maturity that they can realistically be thought of as a platform upon which many applications may be built, including expressive type systems.

Our type-checker, like all static analyzers, has the potential to generate false negatives, that is, rejecting programs as type incorrect that are, in fact, type correct. As any SMT solver is incomplete for the first-order theories that we are interested in, it is possible that the solver is unable to determine an answer to a logical statement. SAGE (Flanagan, 2006; Knowles *et al.*, 2007) avoids these problems by catching these cases and inserting a cast so that the test is performed again at run-time. This has a pleasant effect of not penalizing the developer for any possible incompleteness of an SMT solver. The techniques used in SAGE should apply to Dminor without any great difficulty.

We leave as future work the project of adding support for first-class functions; one direction is to generalize the mixture of syntactic and semantic subtyping introduced by Calcagno *et al.*, (2005).

Finally, the implications of this work go beyond the core calculus Dminor. PADS, JSON, and M, for example, show the significance of programming languages for first-order data. Our work establishes the usefulness of combining refinement types and dynamic type-tests when programming with first-order data, and the viability of statically type-checking such programs with the aid of an SMT solver.

Appendix A Mechanized definition of the intended model

In our logical semantics from Section 3 and its optimized version from Section 6.1, the semantics of a Dminor type is a FOL formula that is interpreted in a specific model. In this section we present the formal definition of this model in the Coq proof assistant (TypiCal Project, 2009). Sorts are encoded as Coq types, and function symbols are interpreted as Coq functions. We focus only on the interpretation of the types and function symbols used by the optimized logical semantics from Section 6.1.

The formalization of the intended Dminor model is valuable for two main reasons. First, the formalization ensures that the model is properly defined. The recursive functions in the model are checked by Coq to be total and terminating. In addition, we have proved in Coq that all functions preserve the logical invariants of the types on which they operate. For instance the FOL sort `Value` is interpreted as the Coq subset (Sozeau, 2006) type $\{x : \text{RawValue} \mid \text{Normal } x\}$, which required us to prove that all functions in the model only produce values in normal form (from values in normal form).

Second, having a natural model that is defined independently of the axioms that are fed to the SMT solver (see Appendix B) allows us to reason about the soundness of these axioms with respect to the model. The soundness of the axioms is not always obvious, since the axioms are meant to have good performance in the

SMT solver rather than to be easy to trust by a human. Proving soundness in a standard model of axioms provided to an automated prover is not new, of course; it was done for instance by Böhme *et al.*, (2008), who proved in Isabelle/HOL that the axiomatization underlying their verification method for C code is correct. The Boyer–Moore family of theorem provers (Boyer *et al.*, 1995) allows both writing logical definitions for models and executing them efficiently (Greve *et al.*, 2008).

Appendix A.1 Values

We first define scalars (sort `General`) and “raw” values as inductive types. Entities are represented as lists of string-raw-value pairs, while collections are represented as lists of raw values.

Model: Raw values

Inductive `General` : Type :=
 | `G_Integer` : `Z` → `General`
 | `G_Text` : `string` → `General`
 | `G_Logical` : `bool` → `General`
 | `G_Null` : `General`.

Inductive `RawValue` : Type :=
 | `G` : `General` → `RawValue`
 | `E` : `list (string * RawValue)` → `RawValue`
 | `C` : `list RawValue` → `RawValue`.

This representation is not canonical, that is, multiple representations for the same value exist, which means we cannot use syntactic equality to compare raw values.

Instead of working directly with raw values, we only consider raw values that are in a normal form. Entities in normal form are sorted by their field name (a string), and do not contain duplicate field names. Collections in normal form are sorted with respect to a total order on raw values (this order is arbitrary but fixed; this order is irrelevant for the semantics of *pure* expressions). The main advantage of using values in normal form is that FOL equality can be interpreted as syntactic equality, as is usual for FOL models.⁷

Model: Sorted string-value maps and value bags

Definition `leAll` (`x` : `A`) (`ys` : `list A`) := `forall y, In y ys → le x y`.

Inductive `Sorted` : `list A` → `Prop` :=
 | `Sorted_nil` : `Sorted nil`
 | `Sorted_cons` : `forall hd tl, leAll hd tl → Sorted tl → Sorted (hd :: tl)`.

⁷ If we had gone with a more complicated interpretation of equality, we would have needed to restrict the interpretation of function symbols to equality-respecting functions, since the interpretation of equality needs to be a congruence.

Definition le_sv ($sv1\ sv2 : (string * RawValue)$) : Prop :=
match $sv1, sv2$ **with**
 ($s1, -$), ($s2, -$) \Rightarrow $cmp_str\ s1\ s2 = Lt \vee cmp_str\ s1\ s2 = Eq$
end.

Definition $Sorted_svm$ ($svm : list (string * RawValue)$) : Prop :=
 $Sorted\ le_sv\ svm$.

Definition le_rval ($v1\ v2 : RawValue$) : Prop :=
 $cmp_rval\ v1\ v2 = Lt \vee cmp_rval\ v1\ v2 = Eq$.

Definition $Sorted_vb$ ($vb : list RawValue$) : Prop := $Sorted\ le_rval\ vb$.

Model: Normal values

Inductive $Normal : RawValue \rightarrow Prop$:=
 | $normal_G : forall\ g, Normal\ (G\ g)$
 | $normal_E : forall\ svm,$
 $NoDup\ (fst\ (split\ svm)) \rightarrow$
 $Sorted_svm\ svm \rightarrow$
 $IndAll\ Normal\ (snd\ (split\ svm)) \rightarrow$
 $Normal\ (E\ svm)$
 | $normal_C : forall\ vb,$
 $Sorted_vb\ vb \rightarrow IndAll\ Normal\ vb \rightarrow Normal\ (C\ vb)$.

We define the Coq type $Value$ (the interpretation of the FOL sort $Value$) as the subset (Sozeau, 2006) of $RawValue$ for which the $Normal$ predicate holds. The sorts $SVMap$ and $VBag$ are interpreted by similar Coq subset types. The elements of $SVMap$ are association lists, lists of key/value pairs; the three conditions in the definition of $SVMap$ require that the list of keys contain no duplicates (so that it forms a finite map), the list of keys is ordered (so that it is in normal form), and each of the contained values is itself in normal form.

Model: Coq types interpreting FOL sorts

Definition $Value := \{x : RawValue \mid Normal\ x\}$.

Definition $SVMap :=$
 $\{svm : list (string * RawValue) \mid NoDup\ (fst\ (split\ svm))$
 $\wedge Sorted_svm\ svm \wedge IndAll\ Normal\ (snd\ (split\ svm))\}$.

Definition $VBag :=$
 $\{vb : list RawValue \mid Sorted_vb\ vb \wedge IndAll\ Normal\ vb\}$.

We define testers and accessors for values and scalars in the usual way. For instance, is_E checks whether its argument is an entity, and if this is the case out_E can be used to obtain the association list corresponding to the entity.

Model: Testers and accessors

```

Program Definition is_E (v : Value) : bool :=
  match v with | E _ => true | _ => false end.
Program Definition out_E (v : Value) : SVMMap :=
  match v with | E svm => svm | _ => nil end.

```

When the argument to `out_E` is not an entity the function returns the empty association list. This choice is arbitrary, but it is necessary that all the functions in the model are total. The testers and accessors for the other constructors are defined in the same way, and for the sake of brevity we omit their definitions here. For collections and entities we further wrap the testers to permit more flexibility in the way they are axiomatized.

Model: Good entities and collections

```

Definition Good_C := is_C.
Definition Good_E := is_E.

```

Appendix A.2 Operations on simple values

The functions `In_Logical`, `In_Integer`, and `In_Text` test whether a value is in the corresponding scalar type.

Model: Testers for simple values

```

Definition In_Logical v := (is_G v) && is_G_Logical (out_G v).
Definition In_Integer v := (is_G v) && is_G_Integer (out_G v).
Definition In_Text v := (is_G v) && is_G_Text (out_G v).

```

We also define shorthand notation for constructing simple values.

Model: Constructors for simple values

```

Program Definition v_tt : Value := G (G_Logical true).
Program Definition v_ff : Value := G (G_Logical false).
Program Definition v_logical (b : bool) : Value := G (G_Logical b).
Program Definition v_int i : Value := G (G_Integer i).
Program Definition v_text s : Value := G (G_Text s).
Program Definition v_null : Value := G (G_Null).

```

The operations on simple values are straightforward and are implemented directly by their Coq counterparts.

Model: Operators on simple values

Definition `O_Sum v1 v2 :=`

`v_int (Zplus (of_G_Integer(out_G v1)) (of_G_Integer(out_G v2))).`

Definition `O_Minus v1 v2 :=`

`v_int (Zminus (of_G_Integer(out_G v1)) (of_G_Integer(out_G v2))).`

Definition `O_Mult v1 v2 :=`

`v_int (Zmult (of_G_Integer(out_G v1)) (of_G_Integer(out_G v2))).`

Definition `O_GT v1 v2 :=`

`match Zcompare (of_G_Integer(out_G v1)) (of_G_Integer(out_G v2))
with Gt => v_tt | _ => v_ff end.`

Definition `O_LT v1 v2 :=`

`match Zcompare (of_G_Integer(out_G v1)) (of_G_Integer(out_G v2))
with Lt => v_tt | _ => v_ff end.`

Definition `O_EQ v1 v2 := v_logical (syn_beq_val v1 v2).`

Definition `O_Not v := v_logical (negb (of_G_Logical (out_G v))).`

Definition `O_And v1 v2 := v_logical (andb (of_G_Logical (out_G v1))
(of_G_Logical (out_G v2))).`

Definition `O_Or v1 v2 := v_logical (orb (of_G_Logical (out_G v1))
(of_G_Logical (out_G v2))).`

Appendix A.3 Operations on entities

The model provides two operations for creating entities: `v_empty` creates an empty entity, and `v_eupdate` creates a new entity from an existing one by updating one field. If the updated field already exists in the original entity then the value of this field will be lost in the new entity. The implementation of `v_eupdate` uses an auxiliary function `update_in_sorted_svm` that implements insertion sorting for association lists. If the key to be added is, however, already present in the association list, then `update_in_sorted_svm` additionally removes the old entry. The `v_eupdate` operation does not correspond to any `Dminor` construct (although it would be easy to add functional entity updates to `Dminor`) but it allows us to construct entity values in an abstract way (without caring how they are implemented – for example, lists versus arrays).

Model: Creating entities

Program Definition `v_empty : Value := E nil.`

Program Definition `v_eupdate (s : string) (v e : Value) : Value :=
E (update_in_sorted_svm (s, v) (out_E e)).`

The two basic operations on entities are: `v_has_field` that tests whether an entity has a certain field, and `v_dot` that given an entity that has a certain field selects the value of this field. Functions `v_has_field` and `v_dot` use the Coq library function `TheoryList.assoc` to obtain the value associated with a given key in a list of pairs.

Model: Basic operations on entities

Program Definition v_has_field ($s : string$) ($v : Value$) : $bool :=$
 $match$ TheoryList.assoc eq_str_dec s (out_E v) **with**
 | Some $v \Rightarrow true$ | None $\Rightarrow false$ **end.**

Program Definition v_dot ($s : string$) ($v : Value$) : $Value :=$
 $match$ TheoryList.assoc eq_str_dec s (out_E v) **with**
 | Some $v \Rightarrow v$ | None $\Rightarrow v_null$ **end.**

Appendix A.4 Operations on collections

The constant v_zero represents the empty collection. The boolean function v_mem tests whether a value is present in a collection using the TheoryList.mem function from the Coq standard library. The function v_add adds an element to a collection using an auxiliary function insert_in_sorted_vb, which implements insertion sorting for collections. In turn v_add is used to define v_add_many , which adds i instances of a given value to a collection.

Model: Functions and predicates on collection

Program Definition v_zero : $Value := C nil.$

Program Definition v_mem ($v cv : Value$) : $bool :=$
 TheoryList.mem eq_rval_dec v (out_C cv).

Program Definition v_add ($v cv : Value$) : $Value :=$
 (C (insert_in_sorted_vb v (out_C cv))).

Fixpoint v_add_many' ($v : Value$) ($n : nat$) ($cv : Value$) : $Value :=$
 $match$ n **with** 0 $\Rightarrow cv$ | S $n' \Rightarrow v_add_many'$ v n' (v_add v cv) **end.**

Definition v_add_many ($v : Value$) ($i : Z$) ($cv : Value$) : $Value :=$
 v_add_many' v (Zabs_nat i) cv.

Definition Closure2 := $Value \rightarrow Value \rightarrow Value.$

Definition v_apply2 ($c : Closure2$) $v1 v2 := c v1 v2.$

Program Fixpoint v_acc_fold ($f : Closure2$) ($vb : VBag$) ($a : Value$)
 {measure List.length vb} : $Value :=$
 $match$ vb **with** nil $\Rightarrow a$ | v :: vb' $\Rightarrow v_acc_fold$ vb' (f a v) **end.**

Definition $v_accumulate$ (clos:Closure2) c := v_acc_fold clos (out_C c).

Finally, $v_accumulate$ implements folding over the elements of a collection using the fixed order on raw values. For the semantics of pure expressions, the order cannot influence the final result.

Appendix B Axiomatization of the Dminor model

We axiomatize the model of Dminor in sorted FOL extended with the background theories of equality, integer arithmetic, algebraic datatypes, and extensional arrays. We only axiomatize the parts of the model that are relevant for the optimized logical semantics in Section 6.1.

In the following we report all the relevant parts of this axiomatization, directly imported from our implementation (file DminorFoundationSmtLib.smt in the

Dminor release). We use the standard SMT-LIB 1.2 format (Ranise & Tinelli, 2006) supported by all recent SMT solvers, together with Z3-specific (de Moura & Bjørner, 2008) extensions for algebraic datatypes and arrays (de Moura & Bjørner, 2009).

We leave it as future work to prove formally that these axioms are properties of the model from Appendix A.

Appendix B.1 An overview on Z3 arrays

We use arrays in our axiomatization to represent collections (multisets) and entities (maps). An array is a function with finite support from one sort (the domain) into another (the range). The domain can be infinite, but the array can differ from a default element only on a finite subset of the domain. As a simple illustrative example (which is not part of our axiomatization for Dminor), we can define an array from integers to Booleans, which basically represent a set of integers using its characteristic function.

Defining an array sort representing sets of integers in Z3

```
:define_sorts ((IntSetArray (Array Int bool)))
```

The basic theory of arrays was introduced by McCarthy (1962) and characterizes functions store and select, using the following two axioms:

$$\begin{aligned} \forall a, i, v. \text{select}(\text{store}(a, i, v), i) &= v \\ \forall a, i, j, v. i = j \vee \text{select}(\text{store}(a, i, v), j) &= \text{select}(a, j) \end{aligned}$$

These axioms can be written in Z3 syntax (for our array sort IntSetArray above) as follows:

The basic theory of arrays in Z3 syntax

```
:assumption (forall (a IntSetArray) (i Int) (v bool)
  (= (select (store a i v) i) v))
:assumption (forall (a IntSetArray) (i Int) (j Int) (v bool)
  (or (= i j) (= (select (store a i v) j) (select a j))))
```

One additional property that is often desirable is extensionality that two arrays are equal when they agree on all elements.

Extensionality of arrays in Z3 syntax

```
:assumption (forall (a1 IntSetArray) (a2 IntSetArray)
  (implies (forall (i Int) (= (select a1 i) (select a2 i))) (= a1 a2)))
```

The select and store function symbols can, for instance, be used to implement a predicate set.contains that checks whether an integer is an element of a set, and a function set.remove to remove an element from a set.

Set membership and removing an element from a set

```

:extrafuns ((set_contains IntSetArray Int bool)
  (set_remove IntSetArray Int IntSetArray))
:assumption (forall (a IntSetArray) (i Int)
  (= (set_contains a i) (select a i)) :pat{ (set_contains a i) })
:assumption (forall (a IntSetArray) (i Int)
  (= (set_remove a i) (store a i false)) :pat{ (set_remove a i) })

```

In the axioms above we use quantifier patterns (de Moura & Bjørner, 2007; Leino & Monahan, 2009) to restrict the number of quantifier instantiations. So the SMT solver will replace `set_remove` by a `store`, but not the other way around. Such careful fine-tuning allows one to choose the right trade-off between performance and completeness. More general quantifier patterns lead to more instantiations of the axioms, which can be expensive and can lead to non-termination. On the other hand, too specific patterns can prevent the SMT solver from even trying to prove useful proof obligations.

Instead of relying directly on the axioms above, Z3 provides an efficient saturation procedure for the extensional array theory as well as a powerful extension called combinatory array logic (de Moura & Bjørner, 2009). The extension defines three new combinators, `const`, `default`, and `map[f]`, which satisfy the following axioms.

Combinatory array logic operations in Z3 syntax

```

:assumption (forall (i Int) (v bool)
  (= (select (const[IntSetArray] v) i) v))
:assumption (forall (v Bool) (= (default (const[IntSetArray] v) v) v))
:assumption (forall (a IntSetArray) (i Int) (v bool)
  (= (default (store a i v) (default a)) v))
:extrafuns ((f bool bool bool))
:assumption (forall (a1 IntSetArray) (a2 IntSetArray) (i Int)
  (= (select (map[f] a1 a2) i) (f (select a1 i) (select a2 i))))

```

For example, we can use these new combinators for defining a constant `set_empty` representing an empty set, a predicate `set_finite` capturing the finiteness of sets, and a function `set_union` that computes the union of two sets.

Additional operations on sets

```

:extrafuns ((set_empty IntSetArray) (set_finite IntSetArray bool)
  (set_union IntSetArray IntSetArray IntSetArray))
:assumption (set_empty = const[IntSetArray] false)
:assumption (forall (a IntSetArray)
  (equiv (set_finite a) (= (default a) false)) :pat{ (set_finite a) })
:assumption (forall (a1 IntSetArray) (a2 IntSetArray)
  (= (set_union a1 a2) (map[and] a1 a2)) :pat{ (set_union a1 a2) })

```

In Z3 all the array axioms above are built-in (so they should not be added manually) and are efficiently implemented (de Moura & Bjørner, 2009).

Appendix B.2 Values

We begin our axiomatization of `Dminor` by defining simple values. For strings and the labels of entities we define a new sort named `String`. The semantics of sorted FOL ensures that this sort is non-empty and disjoint from all other sorts. Since strings and labels are constants and we have no operation on them, we do not further constrain this sort.

The sort `General` is defined as an algebraic datatype with four constructors: `G.Integer` taking a (built-in) integer as argument, `G.Text` taking a `String`, `G.Logical` taking a (built-in) boolean, and the constant `G.Null`.

Simple values

```

:extrasorts (String)
:datatypes ((General
  (G.Integer (of_G.Integer Int))
  (G.Text (of_G.Text String))
  (G.Logical (of_G.Logical bool))
  G.Null))

```

This declaration implicitly defines three accessor functions, named `of_G.Integer`, `of_G.Text`, and `of_G.Logical`, which are inverses to `G.Integer`, `G.Text`, and `G.Logical`. Given an argument of the form `(G.Integer i)`, the function `of_G.Integer` returns `i`; `of_G.Text` and `of_G.Logical` act similarly. In addition, the declaration implicitly defines tester functions by adding the `is_` prefix to the names of each constructor, so `(is_G.Integer g)` tests whether `g` is of the form `(G.Integer i)`, and similarly for `is_G.Text`, `is_G.Logical` and `is_G.Null`.

Values are also defined as a datatype. We use extensional arrays to represent entities and collections. However, since `Z3` syntactically restricts arrays from appearing inside datatypes, and since we need to restrict the arrays so that they represent only finite maps and bags, we use two new (abstract) sorts `SVMap` and `VBag` instead. The sort `SVMap` is then constrained to be isomorphic to the arrays from `Strings` to `Values` for which a finiteness condition holds, while `VBag` is required to be isomorphic to the arrays from `Values` to non-negative `Ints`, again with an additional finiteness condition.⁸

Values

```

:extrasorts (SVMap VBag)
:datatypes ((Value
  (G (out_G General)) ;; simple value (scalar)
  (E (out_E SVMap)) ;; entity: finite map from String to Value
  (C (out_C VBag)))) ;; collection: finite multiset of Value

```

⁸ Because of this additional indirection, our axiomatization of sort `Value` captures not only the values of `Dminor` but also infinite values that contain themselves (for example, a collection that has itself as an element). This is sound, since if a property can be proved of this larger set of values then it also holds for the actual values. In practice it happens very rarely that the SMT solver manages to falsify a property by constructing a cyclic value; still, our code to process a `Z3` model and extract a counterexample (see Section 7) keeps track of cycles and aborts if one is encountered.

Since arrays can in general be infinite, we further restrict the set of values to contain only finite collections and entities using the predicates `Good_C` and `Good_E` (defined below).

Good values

```

:assumption (forall (v Value)
  (implies (Good v)
    (and (implies (is_C v) (Good_C v))
      (implies (is_E v) (Good_E v))))))
:pat { (Good v) }
```

Appendix B.3 Operations on simple values

We define several functions that test whether a value is a Boolean (`In_Logical`), an integer (`In_Integer`), or a string (`In_Text`). These functions are trivial to implement because Z3 already provides testers for datatypes.

Testers for simple values

```

:assumption (forall (v Value)
  (iff (In_Logical v) (and (is_G v) (is_G_Logical (out_G v))))
  :pat { (In_Logical v) })
:assumption (forall (v Value)
  (iff (In_Integer v) (and (is_G v) (is_G_Integer (out_G v))))
  :pat { (In_Integer v) })
:assumption (forall (v Value)
  (iff (In_Text v) (and (is_G v) (is_G_Text (out_G v))))
  :pat { (In_Text v) })
```

We also define more convenient constructors for simple values.

Constructors for simple values

```

:assumption (= v_tt (G(G_Logical true)))
:assumption (= v_ff (G(G_Logical false)))
:assumption (forall (b bool) (= (v_logical b) (G(G_Logical b))))
:pat { (v_logical b) }
:assumption (= v_null (G(G_Null)))
:assumption (forall (n Int) (= (v_int n) (G(G_Integer n))))
:pat { (v_int n) } :pat { (G(G_Integer n)) }
:assumption (forall (s String) (= (v_text s) (G(G_Text s))))
:pat { (v_text s) } :pat { (G(G_Text s)) }
```

The operators on integers and Booleans are easy to define using the built-in SMT-LIB functions.

Operators on simple values

```

:assumption (forall (i1 Int) (i2 Int)
  (= (O_Sum (v_int i1) (v_int i2)) (v_int (+ i1 i2))))
:pat { (O_Sum (v_int i1) (v_int i2)) }
:assumption (forall (v1 Value) (v2 Value)
  (= (O_EQ v1 v2) (ite (= v1 v2) v_tt v_ff)))
:pat { (O_EQ v1 v2) }
:assumption (forall (v Value)
  (= (O_Not v) (ite (not (= v v_tt)) v_tt v_ff)))
:pat { (O_Not v) }
:assumption (forall (v1 Value) (v2 Value)
  (= (O_And v1 v2) (ite (and (= v1 v_tt) (= v2 v_tt)) v_tt v_ff)))
:pat { (O_And v1 v2) }
:assumption (forall (v1 Value) (v2 Value)
  (= (O_Or v1 v2) (ite (or (= v1 v_tt) (= v2 v_tt)) v_tt v_ff)))
:pat { (O_Or v1 v2) }

```

We omit the definitions for `O_NE`, `O_Minus`, `O_Mult`, `O_GT`, and `O_LT`, which follow the same pattern.

Appendix B.4 Operations on entities

Entities

```

:datatypes ((ValueOption
  NoValue
  (SomeValue (of_SomeValue Value))))
:define_sorts ((SVMArray (Array String ValueOption)))
:extrafuns ((alphan SVMArray SVMArray)
  (betam SVMArray SVMArray))

```

We represent entities as arrays from strings to the datatype `ValueOption`, which contains `Values` as well as a special `NoValue` marker. We call such an array `finite` (`FiniteE`) if it has `NoValue` as the default element. We use the functions `alphan` and `betam` as the witnesses of isomorphism between the abstract sort `SVMArray` and the finite part of the array sort `SVMArray`. The axiomatization of entities uses these witness functions intensively.

Operations on entities

```

;; SVMMap and the finite arrays in SVMMapArray are isomorphic
:assumption (forall (am SVMMapArray)
  (implies (FiniteE am) (= (alpham (betam am)) am)))
:assumption (forall (svm SVMMap)
  (and (FiniteE (alpham svm)) (= (betam (alpham svm)) svm)))

:assumption (forall (svm SVMMapArray) (iff (FiniteE svm)
  (= (default svm) NoValue))) :pat{ (FiniteE svm) }

:assumption (forall (v Value)
  (iff (Good_E v) (and (is_E v) (FiniteE (alpham (out_E v))))))
:pat{ (Good_E v) }

:assumption (= v_empty (E (betam (const[SVMMapArray] NoValue))))

:assumption (forall (l String) (v Value) (svm SVMMap)
  (= (v_eupdate l v (E svm))
    (E (betam (store (alpham svm) l (SomeValue v))))))
:pat{ (v_eupdate l v (E svm)) }

:assumption (forall (l String) (svm SVMMap)
  (iff (v_has_field l (E svm)) (not(= (select (alpham svm) l) NoValue))))
:pat { (v_has_field l (E svm)) } ;; :pat (select (alpham svm) l)

:assumption (forall (l String) (svm SVMMap)
  (= (v_dot l (E svm)) (of_SomeValue (select (alpham svm) l))))
:pat { (v_dot l (E svm)) } ;; :pat (select (alpham svm) l)

```

Appendix B.5 Operations on collections

Collections

```

:define_sorts ((VBagArray (Array Value Int)))
:extrafuns ((alphab VBag VBagArray)
  (betab VBagArray VBag))

```

We represent collections as arrays from Values to integers. We call such a collection good (Good.C) when it is finite (default value of the array is zero) and the cardinality of the elements is non-negative. Good collections correspond to finite multisets over values. We use the functions `alphab` and `betab` to define an isomorphism between the abstract sort `VBag` and the good collections in `VBagArray`.

Constraints on bags

```

;; VBag and the finite and positive arrays in VBagArray are isomorphic
:assumption (forall (ab VBagArray)
  (implies (and (Finite ab) (Positive ab)) (= (alphab (betab ab)) ab))
  :pat{ (alphab (betab ab)) })
:assumption (forall (vb VBag)
  (and (Finite (alphab vb)) (Positive (alphab vb))
    (= (betab (alphab vb)) vb))
  :pat{ (betab (alphab vb)) })

;; Good collections are finite and positive
:assumption (forall (v Value)
  (iff (Good_C v)
    (and (is_C v)
      (Finite (alphab (out_C v)))
      (Positive (alphab (out_C v))))))
  :pat{ (Good_C v) })

;; Finiteness of bags
:assumption (forall (a VBagArray)
  (iff (Finite a) (= (default a) 0))
  :pat{ (Finite a) })

;; Only positive indices in bags
:assumption (forall (a VBagArray)
  (iff (Positive a) (forall (v Value) (>= (select a v) 0))
  :pat{ (select a v) }) :pat{ (Positive a) })

```

Closures

```

:extrasorts (Closure2)
:extrafuns ((v_apply2 Closure2 Value Value Value))

```

For axiomatizing `v_accumulate` we use an abstract sort for closures of two arguments (`Closure2`). The `v_apply2` operation defines how each of the closures behaves on its arguments. When giving semantics to an expression **from** x **in** e_1 **let** $y = e_2$ **accumulate** e_3 , a fresh closure is generated for e_2 (called f_{e_2} below), and the following axiom is added for it:

```

:extrafuns ((f_e2 Closure2))
:assumption (forall (x y Value)
  (= (apply2 f_e2 x y)  $\mathbb{V}[[e_2]]$ ) :pat{(apply2 f_e2 x y)})

```

The `v_apply2` function is used in the second axiom for `v_accumulate`.

Operations on collections

```

:assumption (= v_zero (C (betab (const[VBagArray] 0))))

:assumption (forall (v Value) (vb VBag)
  (iff (v_mem v (C vb)) (> (select (alphab vb) v) 0))
  :pat { (v_mem v (C vb)) }
  :pat{ (select (alphab vb) v) } )

:assumption (forall (v Value) (i Int) (vb VBag)
  (= (v_add_many v i (C vb))
    (C (betab (store (alphab vb) v (+ i (select (alphab vb) v))))))
  :pat{ (v_add_many v i (C vb)) } )

:assumption (forall (v Value) (vs Value)
  (= (v_add v vs) (v_add_many v 1 vs)) :pat{ (v_add v vs) } )

;; v_accumulate iterates using an order-preserving function
:assumption (forall (clos Closure2) (initial Value)
  (= (v_accumulate clos v_zero initial) initial)
  :pat{ (v_accumulate clos v_zero initial) } )

:assumption (forall (clos Closure2) (initial Value) (v Value) (vs Value)
  (= (v_accumulate clos (v_add v vs) initial)
    (v_accumulate clos vs (v_apply2 clos v initial)))
  :pat { (v_accumulate clos (v_add v vs) initial) } )

```

Appendix C Proofs

It is immediate from the definition of purity in Section 2.3 that purity is preserved by small-step reduction. This property is used in the proof of Theorem 1.

Lemma 4 (Reduction Preserves Purity)

If e is pure and $e \rightarrow e'$ then e' is pure.

Appendix C.1 Relating operational and logical semantics

In this section we develop proofs for Lemma 2 and Theorem 1 (Full Abstraction) from Section 3.

We begin with a direct inductive definition of the relation $e \Downarrow r$; that is, an error-tracking big-step operational semantics.

Evaluation semantics: $e \Downarrow r$

(Eval Const)

$c \Downarrow \mathbf{Return}(c)$

(Eval Operator 1)

$e_i \Downarrow \mathbf{Return}(v_i) \quad \forall i \in 1..j-1 \quad e_j \Downarrow \mathbf{Error} \quad j \in 1..n$

$\oplus (e_1, \dots, e_n) \Downarrow \mathbf{Error}$

(Eval Operator 2)

$$\frac{e_i \Downarrow \mathbf{Return}(v_i) \quad i \in 1..n \quad \neg \exists v. (\oplus(v_1, \dots, v_n) \mapsto v)}{\oplus(e_1, \dots, e_n) \Downarrow \mathbf{Error}}$$

(Eval Operator 3)

$$\frac{e_i \Downarrow \mathbf{Return}(v_i) \quad \forall i \in 1..n \quad \oplus(v_1, \dots, v_n) \mapsto v}{\oplus(e_1, \dots, e_n) \Downarrow \mathbf{Return}(v)}$$

(Eval Cond 1)

$$\frac{e_1 \Downarrow r \quad r \notin \{\mathbf{Return}(\mathbf{true}), \mathbf{Return}(\mathbf{false})\}}{e_1 ? e_{\mathbf{true}} : e_{\mathbf{false}} \Downarrow \mathbf{Error}}$$

(Eval Cond 2)

$$\frac{e_1 \Downarrow \mathbf{Return}(b) \quad b \in \{\mathbf{true}, \mathbf{false}\} \quad e_b \Downarrow r}{e_1 ? e_{\mathbf{true}} : e_{\mathbf{false}} \Downarrow r}$$

(Eval Let 1)

$$\frac{e_1 \Downarrow \mathbf{Error}}{\mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 \Downarrow \mathbf{Error}}$$

(Eval Let 2)

$$\frac{e_1 \Downarrow \mathbf{Return}(v) \quad e_2\{v/x\} \Downarrow r}{\mathbf{let} \ x = e_1 \ \mathbf{in} \ e_2 \Downarrow r}$$

(Eval Entity 1)

$$\frac{e_i \Downarrow \mathbf{Return}(v_i) \quad \forall i \in 1..j-1 \quad e_j \Downarrow \mathbf{Error} \quad j \in 1..n}{\{\ell_i \Rightarrow e_i \quad i \in 1..n\} \Downarrow \mathbf{Error}}$$

(Eval Entity 2)

$$\frac{e_i \Downarrow \mathbf{Return}(v_i) \quad \forall i \in 1..n}{\{\ell_i \Rightarrow e_i \quad i \in 1..n\} \Downarrow \mathbf{Return}(\{\ell_i \Rightarrow v_i \quad i \in 1..n\})}$$

(Eval Dot 1)

$$\frac{e \Downarrow r \quad \neg \exists v_1, \dots, v_n. (r = \mathbf{Return}(\{\ell_i \Rightarrow v_i \quad i \in 1..n\}) \wedge j \in 1..n)}{e.\ell_j \Downarrow \mathbf{Error}}$$

(Eval Dot 2)

$$\frac{e \Downarrow \mathbf{Return}(\{\ell_i \Rightarrow v_i \quad i \in 1..n\}) \quad j \in 1..n}{e.\ell_j \Downarrow \mathbf{Return}(v_j)}$$

(Eval Collection)

$$\frac{\{v_1, \dots, v_n\} \Downarrow \mathbf{Return}(\{v_1, \dots, v_n\})}{\{v_1, \dots, v_n\} \Downarrow \mathbf{Return}(\{v_1, \dots, v_n\})} \quad \frac{e_1 \Downarrow \mathbf{Error}}{e_1 :: e_2 \Downarrow \mathbf{Error}}$$

(Eval Add 1)

(Eval Add 2)

$$\frac{e_1 \Downarrow \mathbf{Return}(v) \quad e_2 \Downarrow r \quad \neg \exists v_1, \dots, v_n. (r = \mathbf{Return}(\{v_1, \dots, v_n\}))}{e_1 :: e_2 \Downarrow \mathbf{Error}}$$

(Eval Add 3)

$$\frac{e_1 \Downarrow \mathbf{Return}(v) \quad e_2 \Downarrow \mathbf{Return}(\{v_1, \dots, v_n\})}{e_1 :: e_2 \Downarrow \mathbf{Return}(\{v, v_1, \dots, v_n\})}$$

(Eval Appl 1)

$$\frac{e_i \Downarrow \mathbf{Return}(v_i) \quad \forall i \in 1..j-1 \quad e_j \Downarrow \mathbf{Error} \quad j \in 1..n}{f(e_1, \dots, e_n) \Downarrow \mathbf{Error}}$$

(Eval Appl 2)

$$\frac{e_i \Downarrow \mathbf{Return}(v_i) \quad \forall i \in 1..n \quad e\{v_1/x_1\} \dots \{v_n/x_n\} \Downarrow r \quad \text{given function definition } f(x_1 : T_1, \dots, x_n : T_n) : U\{e\}}{f(e_1, \dots, e_n) \Downarrow r}$$

(Eval Accum 1)

$$\frac{e_1 \Downarrow r \quad \neg \exists v_1, \dots, v_n. (r = \mathbf{Return}(\{v_1, \dots, v_n\}))}{\text{from } x \text{ in } e_1 \text{ let } y = e_2 \text{ accumulate } e_3 \Downarrow \mathbf{Error}}$$

(Eval Accum 2)

$$\frac{e_1 \Downarrow \mathbf{Return}(\{v_1, \dots, v_n\}) \quad \text{let } y = e_2 \text{ in let } y = e_3\{v_1/x\} \text{ in } \dots \text{ let } y = e_3\{v_n/x\} \text{ in } y \Downarrow r}{\text{from } x \text{ in } e_1 \text{ let } y = e_2 \text{ accumulate } e_3 \Downarrow r}$$

(Test Wrong) (Test Any)

$$\frac{e \Downarrow \mathbf{Error}}{e \text{ in } T \Downarrow \mathbf{Error}} \quad \frac{e \Downarrow \mathbf{Return}(v)}{e \text{ in Any} \Downarrow \mathbf{Return}(\mathbf{true})}$$

(Test G 1)

$$\frac{e \Downarrow \mathbf{Return}(v) \quad v \in K(G)}{e \text{ in } G \Downarrow \mathbf{Return}(\mathbf{true})}$$

(Test G 2)

$$\frac{e \Downarrow \mathbf{Return}(v) \quad v \notin K(G)}{e \text{ in } G \Downarrow \mathbf{Return}(\mathbf{false})}$$

(Test Entity 1)

$$\frac{e \Downarrow \mathbf{Return}(v) \quad v = \{\ell_i \Rightarrow v_i \mid i \in 1..n\} \wedge j \in 1..n \quad v_j \text{ in } T_j \Downarrow r}{e \text{ in } \{\ell_j : T_j\} \Downarrow r}$$

(Test Entity 2)

$$\frac{e \Downarrow r \quad \neg \exists v_1, \dots, v_n. (r = \mathbf{Return}(\{\ell_i \Rightarrow v_i \mid i \in 1..n\}) \wedge j \in 1..n)}{e \text{ in } \{\ell_j : T_j\} \Downarrow \mathbf{Return}(\mathbf{false})}$$

(Test Collection 1)

$$\frac{e \Downarrow \mathbf{Return}(v) \quad \neg \exists v_1, \dots, v_n. (v = \{v_1, \dots, v_n\})}{e \text{ in } T^* \Downarrow \mathbf{Return}(\mathbf{false})}$$

(Test Collection 2)

$$\frac{e \Downarrow \mathbf{Return}(\{v_1, \dots, v_n\}) \quad v_1 \text{ in } T \ \&\& \dots \ \&\& \ v_n \text{ in } T \Downarrow r}{e \text{ in } T^* \Downarrow r}$$

(Test Refine)

$$\frac{e_1 \Downarrow \mathbf{Return}(v) \quad v \text{ in } T \ \&\& \ e_2\{v/x\} \Downarrow r}{e_1 \text{ in } (x : T \text{ where } e_2) \Downarrow r}$$

Lemma 5

If v is a value then $v \Downarrow \mathbf{Return}(v)$.

Proof

By induction on the structure of v . \square

Lemma 6

Suppose e is closed. If $e \rightarrow e'$ and $e' \Downarrow r$ then $e \Downarrow r$.

Proof

By induction on the derivation of $e' \Downarrow r$, with a case analysis of the reduction $e \rightarrow e'$. We omit the details. \square

Lemma 7

If e is closed and stuck then $e \Downarrow \mathbf{Error}$.

Proof

By induction on the structure of e . We omit the details. \square

By the following lemma, we obtain an independent definition of the relation $e \Downarrow r$ in terms of the reduction relation and stuckness. This is the definition used in Section 2. The equivalent inductive definition given in this section is convenient for proofs.

Lemma 8

Suppose that e is closed.

1. $e \Downarrow \mathbf{Return}(v)$ if and only if $e \rightarrow^* v$.
2. $e \Downarrow \mathbf{Error}$ if and only if there is e' with $e \rightarrow^* e'$ and e' is stuck.

Proof

The forwards direction follow by straightforward inductions on the derivations of $e \Downarrow \mathbf{Return}(v)$ and $e \Downarrow \mathbf{Error}$.

For the reverse direction of (1), we have $e = e_1 \rightarrow \dots \rightarrow e_n \rightarrow v$. By Lemma 5, we have $v \Downarrow \mathbf{Return}(v)$. By repeated applications of Lemma 6, we have $e_i \Downarrow \mathbf{Return}(v)$ for each i from n down to 1, and indeed $e_i \Downarrow \mathbf{Return}(v)$.

For the reverse direction of (2), suppose there is e' such that $e = e_1 \rightarrow \dots \rightarrow e_n = e'$ and e' is stuck. By Lemma 7, we have $e_n \Downarrow \mathbf{Error}$. By repeated applications of Lemma 6, we have $e_i \Downarrow \mathbf{Error}$ for each i from n down to 1, and indeed $e \Downarrow \mathbf{Error}$. \square

Lemma 9

Suppose $\oplus : T_1, \dots, T_n \rightarrow T$.

1. If $\models \mathbf{F}[[T_i]](v_i)$ for each $i \in 1..n$ then there is v such that $\oplus(v_1, \dots, v_n) \mapsto v$.
2. If $\oplus(v_1, \dots, v_n) \mapsto v$ then $\models \mathbf{F}[[T_i]](v_i)$ for each $i \in 1..n$, and $\models \mathbf{F}[[T]](v)$ and $\models \mathbf{O}_{\oplus}(v_1, \dots, v_n) = v$.

The following applies to each operator apart from equality $=$. As mentioned previously, equality is defined on any pair of closed values.

Lemma 10

If $\oplus : G_1, \dots, G_n \rightarrow G$ then $\text{dom}(\oplus) = K(G_1) \times \dots \times K(G_n)$.

Our semantics has the following substitution property.

Lemma 11

1. For all values v and all expressions e that only call labeled-pure functions,

$$\models \mathbf{R}[[e]]\{v/x\} = \mathbf{R}[[e\{v/x\}]]$$

2. For all types T , values v , and FOL terms t :

$$\models \mathbf{F}[\![T]\!](t)\{v/x\} \Leftrightarrow \mathbf{F}[\![T\{v/x\}]\!](t\{v/x\})$$

3. For all types T , values v , and FOL terms t :

$$\models \mathbf{W}[\![T]\!](t)\{v/x\} \Leftrightarrow \mathbf{W}[\![T\{v/x\}]\!](t\{v/x\})$$

Proof

By simultaneous induction on the structure of e and T . \square

We would like to show that if a closed pure expression evaluates a result then that is the result of the expression according to the logical semantics (if $e \Downarrow r$ then $\models \mathbf{R}[\![e]\!] = r$). Intuitively this proof should proceed by induction on the structure of the derivation of $e \Downarrow r$. This works for all cases other than (Eval Accum 2), so it is instructive to observe a failed proof attempt. In this case we know that $e = \mathbf{from} \ x \ \mathbf{in} \ e_1 \ \mathbf{let} \ y = e_2 \ \mathbf{accumulate} \ e_3, \ e_1 \Downarrow \mathbf{Return}(\{v_1, \dots, v_n\})$, and $\mathbf{let} \ y = e_2 \ \mathbf{in} \ \mathbf{let} \ y = e_3\{v_1/x\} \ \mathbf{in} \ \dots \ \mathbf{let} \ y = e_3\{v_n/x\} \ \mathbf{in} \ y \Downarrow r$, for some arbitrary ordering v_1, \dots, v_n . So the induction hypothesis gives us that $\mathbf{R}[\![e_1]\!] = \mathbf{Return}(\{v_1, \dots, v_n\})$ and $\mathbf{Bind} \ y \leftarrow \mathbf{R}[\![e_2]\!] \ \mathbf{in} \ \mathbf{R}[\![\mathbf{let} \ y = e_3\{v_1/x\} \ \mathbf{in} \ \dots \ \mathbf{let} \ y = e_3\{v_n/x\} \ \mathbf{in} \ y]\!] = r$. We need to show that $\mathbf{R}[\![\mathbf{from} \ x \ \mathbf{in} \ e_1 \ \mathbf{let} \ y = e_2 \ \mathbf{accumulate} \ e_3]\!] = r$. By purity we know that for any permutation of v_1, \dots, v_n , including the canonical one used by the model v_{i_1}, \dots, v_{i_n} we have $\mathbf{let} \ y = e_2 \ \mathbf{in} \ \mathbf{let} \ y = e_3\{v_{i_1}/x\} \ \mathbf{in} \ \dots \ \mathbf{let} \ y = e_3\{v_{i_n}/x\} \ \mathbf{in} \ y \Downarrow r$. However, we cannot apply the induction hypothesis to this (possibly) different permutation.

In order to obtain a strong enough induction hypothesis in the accumulate case, we define an auxiliary judgment $e \Downarrow_D r$, which has the same rules as $e \Downarrow r$, with the exception of (Eval Accum 2) that is replaced by the following rule:

Auxiliary evaluation relation: $e \Downarrow_D r$

(Eval Accum D)

$$e_1 \Downarrow_D \mathbf{Return}(\{v_1, \dots, v_n\})$$

$\forall k. v_{i_1}^k, \dots, v_{i_n}^k$ is a permutation of v_1, \dots, v_n

$$\mathbf{let} \ y = e_2 \ \mathbf{in} \ \mathbf{let} \ y = e_3\{v_{i_1}^k/x\} \ \mathbf{in} \ \dots \ \mathbf{let} \ y = e_3\{v_{i_n}^k/x\} \ \mathbf{in} \ y \Downarrow_D r_k$$

$$\mathbf{from} \ x \ \mathbf{in} \ e_1 \ \mathbf{let} \ y = e_2 \ \mathbf{accumulate} \ e_3 \Downarrow_D r_j$$

The new rule (Eval Accum D) does not pick an arbitrary ordering from the start but instead it evaluates using all orderings and only in the end picks one of the results. It is very easy to show that if $e \Downarrow_D r$ then also $e \Downarrow r$.

Lemma 12

If e is closed and $e \Downarrow_D r$ then also $e \Downarrow r$.

Proof

By induction on the structure of the derivation of $e \Downarrow_D r$. \square

If e is additionally pure then also the implication in the other direction holds.

Lemma 13

If e is closed and pure and $e \Downarrow r$ then $e \Downarrow_D r$.

Proof

By induction on the structure of the derivation of $e \Downarrow r$. The proof uses the fact that pure expressions have to terminate on all paths. \square

Lemma 14

For closed and pure e and r , if $e \Downarrow_D r$ then $\models \mathbf{R}[[e]] = r$.

Proof

The proof is by induction on the derivation of $e \Downarrow_D r$. Notice that the purity assumption arises explicitly in the case (Eval Accum D), as well as (Eval Appl 2) for function calls, which also uses Lemma 1. We list these two representative cases of the proof but omit other details.

(Eval Appl 2)

$$\frac{e_i \Downarrow \mathbf{Return}(v_i) \quad \forall i \in 1..n \quad e\{v_1/x_1\} \dots \{v_n/x_n\} \Downarrow r \quad \text{given function definition } f(x_1 : T_1, \dots, x_n : T_n) : U\{e\}}{f(e_1, \dots, e_n) \Downarrow r}$$

By induction hypothesis, we have $\models \mathbf{R}[[e_i]] = \mathbf{Return}(v_i)$ for each $i \in 1..n$. Since $f(e_1, \dots, e_n)$ is pure, it must be that f is a pure-labeled function, and therefore its body e is pure. Hence, by Lemma 1, the definition $f(x_1 : T_1, \dots, x_n : T_n) : U\{e\}$ and e are pure and $e\{v_1/x_1\} \dots \{v_n/x_n\} \Downarrow r$ imply $\models \underline{f}(v_1, \dots, v_n) = r$. We calculate as follows:

$$\begin{aligned} & \models \mathbf{R}[[f(e_1, \dots, e_n)]] \\ &= \mathbf{Bind} \ x_1 \Leftarrow \mathbf{R}[[e_1]] \ \mathbf{in} \ \dots \ \mathbf{Bind} \ x_n \Leftarrow \mathbf{R}[[e_n]] \ \mathbf{in} \\ & \quad \underline{f}(x_1, \dots, x_n) \\ &= \mathbf{Bind} \ x_1 \Leftarrow \mathbf{Return}(v_1) \ \mathbf{in} \ \dots \ \mathbf{Bind} \ x_n \Leftarrow \mathbf{Return}(v_n) \ \mathbf{in} \\ & \quad \underline{f}(x_1, \dots, x_n) \\ &= \underline{f}(v_1, \dots, v_n) \\ &= r \end{aligned}$$

(Eval Accum D)

$$\frac{e_1 \Downarrow_D \mathbf{Return}(\{v_1, \dots, v_n\}) \quad \forall k. v_{i_1^k}, \dots, v_{i_n^k} \text{ is a permutation of } v_1, \dots, v_n \quad \text{let } y = e_2 \ \mathbf{in} \ \text{let } y = e_3\{v_{i_1^k}/x\} \ \mathbf{in} \ \dots \ \text{let } y = e_3\{v_{i_n^k}/x\} \ \mathbf{in} \ y \Downarrow_D r_k}{\text{from } x \ \mathbf{in} \ e_1 \ \text{let } y = e_2 \ \mathbf{accumulate} \ e_3 \Downarrow_D r_j}$$

The induction hypothesis gives us that $\mathbf{R}[[e_1]] = \mathbf{Return}(\{v_1, \dots, v_n\})$ and for any $v_{i_1^k}, \dots, v_{i_n^k}$ permutation of v_1, \dots, v_n $\mathbf{Bind} \ y \Leftarrow \mathbf{R}[[e_2]] \ \mathbf{in} \ \mathbf{R}[[\text{let } y = e_3\{v_{i_1^k}/x\} \ \mathbf{in} \ \dots \ \text{let } y = e_3\{v_{i_n^k}/x\} \ \mathbf{in} \ y]] = r_k$. We choose this to be the canonical permutation used by the model $v_{i_1^k}, \dots, v_{i_n^k}$ and obtain a result r_c for which $\text{let } y = e_2 \ \mathbf{in} \ \text{let } y = e_3\{v_{i_1^k}/x\} \ \mathbf{in} \ \dots \ \text{let } y = e_3\{v_{i_n^k}/x\} \ \mathbf{in} \ y \Downarrow_D r_c$. From Lemma 12 by point (2) in the definition of purity, we obtain $r_c = r_j$. We can thus calculate as follows:

$$\begin{aligned} & \models \mathbf{R}[[\text{from } x \ \mathbf{in} \ e_1 \ \text{let } y = e_2 \ \mathbf{accumulate} \ e_3]] \\ &= \text{res_accumulate}((\mathbf{fun} \ x \ y \rightarrow \mathbf{R}[[e_3]]), v, v') \end{aligned}$$

$$\begin{aligned}
 &= \mathbf{Bind} \ y \leftarrow \mathbf{R}[[e_2]] \ \mathbf{in} \\
 &\quad \mathbf{R}[[\mathbf{let} \ y = e_3\{v_{i_1}^c/x\} \ \mathbf{in} \dots \mathbf{let} \ y = e_3\{v_{i_n}^c/x\} \ \mathbf{in} \ y]]] \\
 &= r_c = r_j
 \end{aligned}$$

The omitted cases proceed similarly. \square

Lemma 15

For closed and pure e and r , if $e \Downarrow r$ then $\models \mathbf{R}[[e]] = r$.

Proof

Immediate from Lemmas 13 and 14. \square

Restatement of Lemma 2

For all closed and pure e and e' , if $e \rightarrow e'$ then $\models \mathbf{R}[[e]] = \mathbf{R}[[e']]$.

Proof

Suppose $e \rightarrow e'$. By Lemma 4, since e is pure, so is e' . By point (2) of the definition of purity, there exists a unique result r such that $e' \Downarrow r$. By Lemma 6, $e \rightarrow e'$ and $e' \Downarrow r$ imply $e \Downarrow r$. By Lemma 15, we have both $\models \mathbf{R}[[e']] = r$ and $\models \mathbf{R}[[e]] = r$. By transitivity, $\models \mathbf{R}[[e]] = \mathbf{R}[[e']]$. \square

Restatement of Theorem 1 (Full Abstraction)

For all closed pure expressions e and e' , we have $\models \mathbf{R}[[e]] = \mathbf{R}[[e']]$ if and only if, for all r , $e \Downarrow r \Leftrightarrow e' \Downarrow r$.

Proof

Since e and e' are closed and pure, by point (2) of the definition of purity there exist unique results r and r' such that $e \Downarrow r$ and $e' \Downarrow r'$. By Lemma 15, we have $\models \mathbf{R}[[e]] = r$ and $\models \mathbf{R}[[e']] = r'$. Given these facts, we have $\models \mathbf{R}[[e]] = \mathbf{R}[[e']]$ if and only if $r = r'$ if and only if for all r'' , $e \Downarrow r'' \Leftrightarrow e' \Downarrow r''$. \square

Appendix C.2 Algorithmic purity implies purity

Here we develop the proof for Theorem 2 from Section 4, which shows that algorithmic purity implies purity. We start with a series of useful lemmas.

Lemma 16

Values are algorithmically pure.

Proof

This is straightforward since none of the restrictions required for algorithmic purity applies to values. \square

Lemma 17

If e is algorithmically pure, so is $e\{v/x\}$ for any v .

Proof

By induction on the structure of e . The variable base case x is handled by Lemma 16. The function application case is straightforward since substitution only applies to function arguments, and not to function bodies. In the accumulate case we have

$$\models \mathbf{R}[\llbracket \text{let } y = e_3\{x_1/x\}\{y_1/y\} \text{ in } e_3\{x_2/x\} \rrbracket] = \mathbf{R}[\llbracket \text{let } y = e_3\{x_2/x\}\{y_1/y\} \text{ in } e_3\{x_1/x\} \rrbracket]$$

and need to show the following, where we may assume that the bound variable y is distinct from x .

$$\begin{aligned} & \models \mathbf{R}[\llbracket \text{let } y = e_3\{v/x\}\{x_1/x\}\{y_1/y\} \text{ in } e_3\{v/x\}\{x_2/x\} \rrbracket] = \\ & \mathbf{R}[\llbracket \text{let } y = e_3\{v/x\}\{x_2/x\}\{y_1/y\} \text{ in } e_3\{v/x\}\{x_1/x\} \rrbracket] \end{aligned}$$

This equation follows from Lemma 11 together with the substitution property of the logic (remember that in FOL free variables are implicitly universally quantified). Condition (3) in the definition of algorithmic purity follows directly by the induction hypothesis. \square

Lemma 18

If e is algorithmically pure, so is $e\sigma$ for any value substitution σ .

Proof

Since e only has finitely many free variables, we only need to consider non-empty finite substitutions of the form $\sigma = \{v_0/x_0\} \dots \{v_n/x_n\}$. The proof proceeds by induction on n , using Lemma 17 in the inductive case. \square

Lemma 19

If e is algorithmically pure and $e \rightarrow e'$ then e' is also algorithmically pure.

Proof

By induction on the derivation of $e \rightarrow e'$, using the following equivalent inductive formulation of algorithmic purity.

Reformulation of algorithmic purity: $\vdash T$ pure $\vdash e$ pure

-
- \vdash Any pure always
 - $\vdash G$ pure always
 - $\vdash T^*$ pure, if $\vdash T$ pure
 - $\vdash \{\ell : T\}$ pure, if $\vdash T$ pure
 - $\vdash (x : T \text{ where } e)$ pure, if $\vdash T$ pure and $\vdash e$ pure
 - $\vdash x$ pure always
 - $\vdash c$ pure always
 - $\vdash \oplus(e_1, \dots, e_n)$ pure, if $\vdash e_i$ pure for each $i \in 1..n$
 - $\vdash e_1 ? e_2 : e_3$ pure, if $\vdash e_i$ pure for each $i \in 1..3$
 - $\vdash \text{let } x = e_1 \text{ in } e_2$ pure, if $\vdash e_1$ pure and $\vdash e_2$ pure
 - $\vdash e \text{ in } T$ pure, if $\vdash e$ pure and $\vdash T$ pure
 - $\vdash \{\ell_i \Rightarrow e_i \mid i \in 1..n\}$ pure, if $\vdash e_i$ pure for each $i \in 1..n$
 - $\vdash e.\ell$ pure if $\vdash e$ pure
 - $\vdash \{v_1, \dots, v_n\}$ pure always
 - $\vdash e_1 :: e_2$ pure if $\vdash e_1$ pure and $\vdash e_2$ pure

\vdash **from** x **in** e_1 **let** $y = e_2$ **accumulate** e_3 **pure**
 if $\vdash e_1$ **pure** and $\vdash e_2$ **pure** and $\vdash e_3$ **pure**
 and $\models \mathbf{R}[\text{let } y = e_3\{x_1/x\}\{y_1/y\} \text{ in } e_3\{x_2/x\}] =$
 $\mathbf{R}[\text{let } y = e_3\{x_2/x\}\{y_1/y\} \text{ in } e_3\{x_1/x\}]$
 (where the variables x_1 , x_2 , and y_1 do not appear free in e_3)
 $\vdash f(e_1, \dots, e_n)$ **pure** if $\vdash e_i$ **pure** for each $i \in 1..n$ and f is labeled-pure

The rest of the proof is routine and was mechanized in Coq. \square

We show that the reduction relation is terminating on algorithmically pure expressions.

Lemma 20

All reduction sequences starting from closed algorithmically pure expressions are finite.

Proof sketch

Recursive functions have to decrease the size of their arguments on each recursive call, which guarantees their termination. The only other source of repetitive computation is accumulate expressions. But collections are finite and the accumulates are immediately inlined, so this will again always terminate. \square

Lemma 21

If e is closed and algorithmically pure then there exists (at least) a result r so that $e \Downarrow r$.

Proof

Immediate from Lemma 20. \square

The most important step for showing the uniqueness of evaluation results for algorithmically pure expressions is to show that the result of evaluating such expressions coincides with the result provided by the logical semantics.

Lemma 22

For all closed and algorithmically pure expressions e and for all results r , if $e \Downarrow r$ then $\models \mathbf{R}[e] = r$.

Proof

By induction on the structure of the derivation of $e \Downarrow r$, with appeal to the big-step semantics in Appendix C.1. \square

Lemma 23

If e is closed and algorithmically pure and $e \Downarrow r_1$, and $e \Downarrow r_2$ then $r_1 = r_2$.

Proof

By Lemma 22 we have $\models \mathbf{R}[e] = r_1$ and $\models \mathbf{R}[e] = r_2$. By transitivity it follows that $\models r_1 = r_2$, which directly implies that $r_1 = r_2$. \square

Lemma 24 (Subexpressions and Substitution)

For all value substitutions σ , if e' is a subexpression of $e\sigma$ then there exists e'' so that e'' is a subexpression of e and $e' = e''\sigma$.

Restatement of Theorem 2

If e is algorithmically pure then e is pure.

Proof

We prove the following more general statement by mutual induction on e and T :

- a. If e is algorithmically pure then e is pure.
- b. For all T , if e is a subexpression of T and e is algorithmically pure then e is pure.

For proving a., by definition, e is pure if and only if for any value substitution σ each of the following four properties hold:

1. $e\sigma$ is terminating,
2. there exists a unique result r such that $e\sigma \Downarrow r$,
3. for every subexpression $f(e_1, \dots, e_n)$ of $e\sigma$, the function f is labeled-pure, and
4. all subexpressions of $e\sigma$ are pure.

Let σ be an arbitrary value substitution. By Lemma 18 we have that $e\sigma$ is algorithmically pure. The first three properties can be proved immediately without using the induction hypothesis. Property (1) follows from Lemma 20. Property (2) follows from Lemmas 21 and 23. Property (3) is immediate from the definition of algorithmic purity.

The only property that uses the induction hypothesis is Property (4): all subexpressions of $e\sigma$ are pure. Let e' be an arbitrary subexpression of $e\sigma$. By Lemma 24 we have that there exists e'' so that e'' is a subexpression of e and $e' = e''\sigma$. We need to prove that $e''\sigma$ is pure, and we do this by case analysis on e . All cases follow immediately by applying the induction hypothesis. The only exception is when $e = e_0$ in T and e'' is a subexpression of T , but there we can use b.. The proof of b. is also simple, by case analysis on the T , and uses the main induction hypothesis when T is a refinement type ($x : T_0$ **where** e_0) and e'' is a subexpression of e_0 . \square

Appendix C.3 Logical soundness

Here we develop the proof for Theorem 3 that relates type assignment to the logical semantics of types and expressions. We start with a series of useful lemmas.

Lemma 25 (Transitivity of Semantic Subtyping)

If $E \vdash T <: T'$ and $E \vdash T' <: T''$ then $E \vdash T <: T''$.

Proof

By expanding definitions. \square

Lemma 26

If e is alg. pure then: $\models \mathbf{F}[\![\text{Ok}(e)]\!](t) \Leftrightarrow (\mathbf{R}[\![e]\!] = \mathbf{Return}(\mathbf{true}))$.

Proof

We have:

$$\begin{aligned}
 & \models \mathbf{F}[\![\text{Ok}(e)]\!](t) \\
 & \Leftrightarrow \mathbf{F}[\![x : \text{Any where } e]\!](t) \quad x \notin \text{fv}(e) \\
 & \Leftrightarrow \mathbf{F}[\![\text{Any}]\!](t) \wedge \text{let } x = t \text{ in } (\mathbf{R}[\![e]\!] = \mathbf{Return}(\text{true})) \\
 & \Leftrightarrow (\mathbf{R}[\![e]\!] = \mathbf{Return}(\text{true}))
 \end{aligned}$$

□

Lemma 27

If e is alg. pure then: $\models \mathbf{F}[\![e : T]\!](t) \Leftrightarrow \mathbf{F}[\![T]\!](t) \wedge (\mathbf{R}[\![e]\!] = \mathbf{Return}(t))$.

Proof

We have:

$$\begin{aligned}
 & \models \mathbf{F}[\![e : T]\!](t) \\
 & \Leftrightarrow \mathbf{F}[\![x : T \text{ where } x == e]\!](t) \quad x \notin \text{fv}(e) \\
 & \Leftrightarrow \mathbf{F}[\![T]\!](t) \wedge \text{let } x = t \text{ in } (\mathbf{R}[\![x == e]\!] = \mathbf{Return}(\text{true})) \\
 & \Leftrightarrow \mathbf{F}[\![T]\!](t) \wedge \\
 & \quad (\mathbf{Bind } y \Leftarrow \mathbf{R}[\![e]\!] \text{ in } \mathbf{Return}(\text{v_logical}(t = y))) = \mathbf{Return}(\text{true}) \\
 & \Leftrightarrow \mathbf{F}[\![T]\!](t) \wedge (\mathbf{R}[\![e]\!] = \mathbf{Return}(t))
 \end{aligned}$$

□

The following lemma characterizes singular subtyping in terms of the logical semantics.

Lemma 28 (Singular Subtyping)

Suppose $E \vdash e : T$ and $E \vdash T'$ and $x \notin \text{dom}(E)$.

1. If e is alg. pure then

$$E \vdash [e : T] <: T' \text{ iff } \models \mathbf{F}[\![E]\!] \wedge \mathbf{F}[\![T]\!](\text{out_v}(\mathbf{R}[\![e]\!])) \Longrightarrow \mathbf{F}[\![T']]\!(\text{out_v}(\mathbf{R}[\![e]\!]))$$

2. If e is not alg. pure then

$$E \vdash [e : T] <: T' \text{ iff } \models \mathbf{F}[\![E]\!] \wedge \mathbf{F}[\![T]\!](x) \Longrightarrow \mathbf{F}[\![T']]\!(x)$$

Proof

In case (1), we have

$$\begin{aligned}
 & E \vdash [e : T] <: T' \\
 & \text{iff } \models (\mathbf{F}[\![E]\!] \wedge \mathbf{F}[\![e : T]\!](x)) \Rightarrow \mathbf{F}[\![T']]\!(x) \\
 & \text{iff } \models (\mathbf{F}[\![E]\!] \wedge \mathbf{F}[\![T]\!](x) \wedge (\mathbf{R}[\![e]\!] = \mathbf{Return}(x))) \Rightarrow \mathbf{F}[\![T']]\!(x) \\
 & \text{iff } \models (\mathbf{F}[\![E]\!] \wedge \mathbf{F}[\![T]\!](x) \wedge (\text{out_v}(\mathbf{R}[\![e]\!]) = x)) \Rightarrow \mathbf{F}[\![T']]\!(x) \\
 & \text{iff } \models (\mathbf{F}[\![E]\!] \wedge \mathbf{F}[\![T]\!](\text{out_v}(\mathbf{R}[\![e]\!])) \Rightarrow \mathbf{F}[\![T']]\!(\text{out_v}(\mathbf{R}[\![e]\!]))
 \end{aligned}$$

In case (2), $[e : T] = T$, and the calculation is immediate. □

By the following lemma, singular subtyping is transitive, and hence we have that any derivation of a type assignment can be seen as one instance of a structural rule

plus one instance of (Exp Singular Subsum). This observation is useful, for example, in proving type preservation, Theorem 4.

Lemma 29 (Transitivity of Singular Subtyping)

If $E \vdash [e : T] <: T'$ and $E \vdash [e : T'] <: T''$ then $E \vdash [e : T] <: T''$.

Proof

An easy application of Lemma 28. □

We can now prove the logical soundness of the type system.

Restatement of Theorem 3 (Logical Soundness)

1. If e is alg. pure and $E \vdash e : T$ then
 - a. $\models \mathbf{F}[[E]] \Rightarrow \text{Proper}(\mathbf{R}[[e]])$
 - b. $\models \mathbf{F}[[E]] \Rightarrow \mathbf{F}[[T]](\text{out}.\mathbf{V}(\mathbf{R}[[e]]))$
2. If $E \vdash U$ then $\models \mathbf{F}[[E]] \Rightarrow \forall y. \neg \mathbf{W}[[U]](y)$, for $y \notin \text{fv}(U)$.

Proof

By mutual induction on the derivation of judgments. A detailed argument appears in the technical report.

□

Appendix C.4 Preservation and progress

Here we develop proofs of Theorems 4 and 5 that imply the safety of our declarative type system from Section 5.

We have the following basic properties.

Lemma 30 (Implied Judgments)

1. If $E \vdash T$ then $E \vdash \diamond$ and $\text{fv}(T) \subseteq \text{dom}(E)$.
2. If $E \vdash T <: T'$ then $E \vdash T$ and $E \vdash T'$.
3. If $E \vdash e : T$ then $E \vdash T$ and $\text{fv}(e) \subseteq \text{dom}(E)$.

Proof

By simultaneous induction on the derivations of each judgment. □

Lemma 31 (All Values Typable)

For any v we have $E \vdash v : \text{Any}$.

Proof

By induction on the structure of v . □

Lemma 32 (Weakening)

Suppose $E, x : T' \vdash \diamond$ and $x \notin \text{dom}(E')$.

1. If $E, E' \vdash \diamond$ then $E, x : T', E' \vdash \diamond$.
2. If $E, E' \vdash T$ then $E, x : T', E' \vdash T$.
3. If $E, E' \vdash S <: T$ then $E, x : T', E' \vdash S <: T$.
4. If $E, E' \vdash e : T$ then $E, x : T', E' \vdash e : T$.

Proof

The proof is by simultaneous induction on the derivation of the judgments $E, E' \vdash \diamond$ and $E, E' \vdash T$ and $E, E' \vdash S <: T$ and $E, E' \vdash e : T$. \square

Lemma 33 (Bound Weakening)

Suppose $E \vdash T <: T'$.

1. If $E, x : T', E' \vdash \diamond$ then $E, x : T, E' \vdash \diamond$.
2. If $E, x : T', E' \vdash S$ then $E, x : T, E' \vdash S$.
3. If $E, x : T', E' \vdash S <: S'$ then $E, x : T, E' \vdash S <: S'$.
4. If $E, x : T', E' \vdash e : S$ then $E, x : T, E' \vdash e : S$.

Proof

By simultaneous induction on derivations. \square

Lemma 34 (Semantic Substitution)

1. For all e', x , alg. pure e so that $E \vdash e : V$ we have that
 $\models \mathbf{F}[\![E]\!] \Rightarrow \mathbf{R}[\![e']\!]\{\text{out.V}(\mathbf{R}[\![e]\!])/x\} = \mathbf{R}[\![e'\{e/x\}]\!]$;
2. For all T, t , alg. pure e so that $E \vdash e : V$ we have that
 - a. $\models \mathbf{F}[\![E]\!] \Rightarrow \mathbf{F}[\![T]\!](t)\{\text{out.V}(\mathbf{R}[\![e]\!])/x\} \Leftrightarrow \mathbf{F}[\![T\{e/x\}]\!](t\{e/x\})$;
 - b. $\models \mathbf{F}[\![E]\!] \Rightarrow \mathbf{W}[\![T]\!](y)\{\text{out.V}(\mathbf{R}[\![e]\!])/x\} \Leftrightarrow \mathbf{W}[\![T\{e/x\}]\!](t\{e/x\})$.

Proof

By mutual induction on the structure of e' and T . A detailed argument appears in the technical report. \square

Lemma 35

For all E, E' , alg. pure e, x , if $E \vdash e : T$ then

$$\models \mathbf{F}[\![E]\!] \Rightarrow \mathbf{F}[\![E']\!]\{\text{out.V}(\mathbf{R}[\![e]\!])/x\} \Leftrightarrow \mathbf{F}[\![E'\{e/x\}]\!]$$

Proof

By induction on the structure of E' , with appeal to Lemma 34. \square

Lemma 36 (Lookup)

If $E \vdash \diamond$ and $(x : T) \in E$ and $(x : T') \in E$ then $T = T'$.

Proof

If $E \vdash \diamond$ all the entries in E are for distinct variables. \square

Lemma 37 (Substitution)

Suppose $E \vdash e' : T'$ and e' alg. pure.

1. If $E, x : T', E' \vdash \diamond$ then $E, E'\{e'/x\} \vdash \diamond$.
2. If $E, x : T', E' \vdash T$ then $E, E'\{e'/x\} \vdash T\{e'/x\}$.
3. If $E, x : T', E' \vdash S <: T$ then $E, E'\{e'/x\} \vdash S\{e'/x\} <: T\{e'/x\}$.
4. If $E, x : T', E' \vdash e : T$ then $E, E'\{e'/x\} \vdash e\{e'/x\} : T\{e'/x\}$.

Proof

The proof is by simultaneous induction on the derivation of the judgments. A detailed argument appears in the technical report. \square

Lemma 38

1. If $\emptyset \vdash e : T$, $\emptyset \vdash e' : T$, $e \rightarrow e'$ and e is alg. pure then $\emptyset \vdash [e' : T] <: [e : T]$.
2. If $\emptyset \vdash e : \text{Logical}$, $\emptyset \vdash e' : \text{Logical}$, $e \rightarrow e'$ then $\emptyset \vdash \text{Ok}(e') <: \text{Ok}(e)$.

Proof

1. From Lemma 19 we have that e' is also alg. pure. We have that $\mathbf{F}[[e' : T]](x)$ holds iff $\mathbf{F}[[T]](x) \wedge (\mathbf{R}[[e']] = \mathbf{Return}(x))$. By Lemma 2 this is equivalent to $\mathbf{F}[[T]](x) \wedge (\mathbf{R}[[e]] = \mathbf{Return}(x))$, which is equivalent to $\mathbf{F}[[e : T]](x)$.
2. We proceed by considering e .

Case e is alg. pure. From Lemma 19 we have that e' is also alg. pure. We have that $\mathbf{F}[[\text{Ok}(e')]](x)$ holds iff $\mathbf{F}[[\text{Any}]](x) \wedge (\mathbf{R}[[e']] = \mathbf{Return}(x))$. By Lemma 2 this is equivalent to $\mathbf{F}[[\text{Any}]](x) \wedge (\mathbf{R}[[e]] = \mathbf{Return}(x))$, which is equivalent to $\mathbf{F}[[\text{Ok}(e)]](x)$.

Case e is not alg. pure. By definition $\text{Ok}(e)$ is Any, and so our property holds as Any is the top type.

□

Another useful lemma relates subtyping with reduction via substitution.

Lemma 39

If $\emptyset \vdash e : U$, $\emptyset \vdash e' : U$, $e \rightarrow e'$ and e is alg. pure then $\emptyset \vdash T\{e/x\} <: T\{e'/x\}$ and $\emptyset \vdash T\{e'/x\} <: T\{e/x\}$.

Proof

Assume $\mathbf{F}[[T\{e/x\}]](t)$ for some term t . Then by Lemma 34, $\mathbf{F}[[T\{e'/x\}]](t)$ holds just if $\mathbf{F}[[T]](t)\{\text{out}.\forall(\mathbf{R}[[e]])/x\}$. By Lemma 2 this is equivalent to $\mathbf{F}[[T]](t)\{\text{out}.\forall(\mathbf{R}[[e']])/x\}$. Again by Lemma 34 this holds if and only if $\mathbf{F}[[T\{e'/x\}]](t)$, as e' is alg. pure by Lemma 19. □

Before we proceed to the preservation theorem, we first need some inversion lemmas for entity and collection types.

Lemma 40 (Entity Type Inversion)

1. If $E \vdash \{\ell_i \Rightarrow v_i\}_{i \in 1..n} : \{\ell_i : T_i\}_{i \in 1..n}$ then $E \vdash v_i : T_i$, for $i \in 1..n$.
2. If $E \vdash \{\ell_i \Rightarrow v_i\}_{i \in 1..n} : \text{Any}$ then $E \vdash v_i : \text{Any}$, for $i \in 1..n$.

Lemma 41 (Collection Type Inversion)

1. If $E \vdash \{v_1, \dots, v_n\} : T^*$ then $E \vdash v_i : T$, for $i \in 1..n$.
2. If $E \vdash \{v_1, \dots, v_n\} : \text{Any}$ then $E \vdash v_i : \text{Any}$, for $i \in 1..n$.

We also need the following lemma, which captures the intuition that if we know that a value inhabits a type, then assuming that it does not inhabit that type leads to a degenerative subtype relation.

Lemma 42

If $E \vdash v : T$ then $E, - : \text{Ok}(!v \text{ in } T) \vdash U <: V$, for any types U, V such that $E \vdash U$ and $E \vdash V$.

Restatement of Theorem 4 (Preservation)

If $\emptyset \vdash e : T$ and $e \rightarrow e'$ then $\emptyset \vdash e' : T$.

Proof

By induction on the derivation of $\emptyset \vdash e : T$. A detailed argument appears in the technical report. \square

Lemma 43 (Canonical Forms)

1. If $\emptyset \vdash v : \text{Integer}$ then $v = \underline{i}$ for some integer i .
2. If $\emptyset \vdash v : \text{Text}$ then $v = \underline{s}$ for some string s .
3. If $\emptyset \vdash v : \text{Logical}$ then $v = \text{true}$ or $v = \text{false}$.
4. If $\emptyset \vdash v : T^*$ then $v = \{v_1, \dots, v_n\}$ for some values v_1, \dots, v_n where $\emptyset \vdash v_i : T$ for all i .
5. If $\emptyset \vdash v : \{\ell : T\}$ then $v = \{\ell_i \Rightarrow v_i \mid i \in 1..n\}$ for some values v_1, \dots, v_n and field names ℓ_1, \dots, ℓ_n so that $\ell = \ell_i$ for some i , and additionally $\emptyset \vdash v_i : T$.

Proof

All parts of the lemma are proved by first applying part (1) (a) of Theorem 3 to the typing derivation of v . For instance for (3) we obtain that $\models \mathbf{F}[\emptyset] \Rightarrow \mathbf{F}[\text{Logical}](\text{out.V}(\text{Return}(v)))$. This is equivalent to $\models \text{In_Logical}(v)$, which gives us that $v = \text{true}$ or $v = \text{false}$ by the definition of In_Logical in the model. The other cases proceed in exactly the same way. \square

Lemma 44 (Progress for Type-tests)

For all types T and values v , $\exists e. v \text{ in } T \rightarrow e$.

Proof

By induction on the structure of T using the reduction rules for type-tests. \square

Restatement of Theorem 5 (Progress)

If $\emptyset \vdash e : T$ and e is not a value then $\exists e'. e \rightarrow e'$.

Proof

By induction on the derivation of $\emptyset \vdash e : T$ using Lemma 44 in the (Exp Test) case. The (Exp Operator), (Exp Cond), (Exp Dot), (Exp Add), and (Exp Acc) cases use Lemma 43. \square

Appendix C.5 Soundness of the algorithmic type system

In this section we present the soundness proof for the algorithmic type system from Section 6.2.

The key property of type normalization is that it preserves the semantics of types. First we state the following properties of the helper functions that are used in type normalization.

Lemma 45 (Soundness of Helper Functions)

1. If $E \vdash R_1$ and $E \vdash R_2$ then $E \vdash R_1 \ \& \ R_2 <: \text{Conj}_{RR}(R_1, R_2)$ and $E \vdash \text{Conj}_{RR}(R_1, R_2) <: R_1 \ \& \ R_2$.
2. If $E \vdash R_1$ and $E \vdash D_1$ then $E \vdash R_1 \ \& \ D_1 <: \text{Conj}_{RD}(R_1, D_1)$ and $E \vdash \text{Conj}_{RD}(R_1, D_1) <: R_1 \ \& \ D_1$.
3. If $E \vdash D_1$ and $E \vdash D_2$ then $E \vdash D_1 \ \& \ D_2 <: \text{Conj}_{DD}(D_1, D_2)$ and $E \vdash \text{Conj}_{DD}(D_1, D_2) <: D_1 \ \& \ D_2$.

Before we proceed to the soundness of type normalization, we state some properties of semantic subtyping that are immediate by definition.

Lemma 46

1. If $E \vdash S$, $E \vdash T$ and $E \vdash S <: T$ then $E \vdash (x : S \text{ where } e) <: (x : T \text{ where } e)$.
2. $E \vdash (x : T_1 \ \& \ T_2 \text{ where } e) <: (x : T_1 \text{ where } e) \ \& \ (x : T_2 \text{ where } e)$.
3. $E \vdash (x_1 : (x_2 : T \text{ where } e_2) \text{ where } e_1) <: (x_1 : T \text{ where } e_1) \ \& \ (x_2 : T \text{ where } e_2)$.

Lemma 47 (Soundness of Type Normalization)

1. If $E \vdash T$ and $norm(T) = D$ then D is a normal type with $E \vdash T <: D$.
2. If $E \vdash (x : C \text{ where } e)$ and $norm_r(x : C \text{ where } e) = D$ then $E \vdash (x : C \text{ where } e) <: D$.

Proof

Proof by mutual induction on T and e . Most of the cases are routine, here we give just two.

Case T is of the form $(x : T \text{ where } e)$. We have that $norm(T) = \prod_{i=1}^n (x_i : C_i \text{ where } e_i)$ and by induction on T , $E \vdash T <: \prod_{i=1}^n (x_i : C_i \text{ where } e_i)$. By part (1) of Lemma 46 we have that $E \vdash (x : T \text{ where } e) <: (x : \prod_{i=1}^n (x_i : C_i \text{ where } e_i) \text{ where } e)$. By part (2) of Lemma 46 we have that $E \vdash (x : T \text{ where } e) <: (\prod_{i=1}^n (x : (x_i : C_i \text{ where } e_i) \text{ where } e))$. By part (3) of Lemma 46 we have that $E \vdash (x : T \text{ where } e) <: (\prod_{i=1}^n ((x_i : C_i \text{ where } e_i) \ \& \ (x : C_i \text{ where } e)))$. We also have by mutual induction $E \vdash (x : C_i \text{ where } e) <: norm_r(x : C_i \text{ where } e)$, so we can deduce that $E \vdash (x : T \text{ where } e) <: (\prod_{i=1}^n ((x_i : C_i \text{ where } e_i) \ \& \ (norm_r(x : C_i \text{ where } e))))$. By Lemma 45 we can conclude $E \vdash (x : T \text{ where } e) <: (\prod_{i=1}^n Conj_{DD}(x_i : C_i \text{ where } e_i, norm_r(x : C_i \text{ where } e)))$ as required.

Case e is of the form $x \text{ in } T$. We have by definition that $norm_r(x : C \text{ where } (x \text{ in } T)) = norm(C \ \& \ T)$. By mutual induction we have that $E \vdash (C \ \& \ T) <: norm(C \ \& \ T)$. We assume $F[[E]]$ and we have that $F[[x : C \text{ where } (x \text{ in } T)]](t)$ is equal to the following by expanding definitions:

$$F[[C]](t) \wedge \text{let } x = t \text{ in } (\mathbf{R}[[x \text{ in } T]] = \mathbf{Return}(\text{true}))$$

By further expansion and Theorem 3 this is equivalent to

$$F[[C]](t) \wedge \text{let } x = t \text{ in} \\ ((\text{if } F[[T]](x) \text{ then } \mathbf{Return}(\text{true}) \text{ else } \mathbf{Return}(\text{false})) = \mathbf{Return}(\text{true})),$$

which is clearly equivalent to $F[[C]](t) \wedge F[[T]](t)$, which by Theorem 3 and the meaning of $E \vdash (C \ \& \ T) <: norm(C \ \& \ T)$ allows us to deduce $F[[norm(C \ \& \ T)]](t)$ as required. \square

Lemma 48 (Soundness of Field Type Extraction)

1. If $E \vdash A$ and $A.\ell \rightsquigarrow U$ then $E \vdash A <: \{\ell : U\}$.
2. If $E \vdash C$ and $C.\ell \rightsquigarrow U$ then $E \vdash C <: \{\ell : U\}$.
3. If $E \vdash R$ and $R.\ell \rightsquigarrow U$ then $E \vdash R <: \{\ell : U\}$.
4. If $E \vdash D$ and $D.\ell \rightsquigarrow U$ then $E \vdash D <: \{\ell : U\}$.

Proof

All parts follow from expanding definitions; here we consider part (3). If $E \vdash R$ and $R.\ell \rightsquigarrow U$ then it must be the case that R is of the form $(x : C \textbf{ where } e)$, and $C.\ell \rightsquigarrow U$. By part (2) we know that $E \vdash C <: \{\ell : U\}$. We know that for any type C with $E \vdash C$ that $E \vdash (x : C \textbf{ where } e) <: C$, and so by transitivity (Lemma 25) we can conclude $E \vdash (x : C \textbf{ where } e) <: \{\ell : U\}$. \square

Lemma 49 (Soundness of Item Type Extraction)

1. If $E \vdash A$ and $A.\text{Items} \rightsquigarrow U$ then $E \vdash A <: U^*$.
2. If $E \vdash C$ and $C.\text{Items} \rightsquigarrow U$ then $E \vdash C <: U^*$.
3. If $E \vdash R$ and $R.\text{Items} \rightsquigarrow U$ then $E \vdash R <: U^*$.
4. If $E \vdash D$ and $D.\text{Items} \rightsquigarrow U$ then $E \vdash D <: U^*$.

Proof

Similar to the proof of the previous lemma; we omit the details. \square

Lemma 50 (Synthesis Checkable)

If $E \vdash e \rightarrow T$ then $E \vdash e \leftarrow T$.

Proof

By (Swap) and reflexivity of singular subtyping. \square

Restatement of Theorem 7 (Soundness of Algorithmic Type System)

1. If $E \triangleright \diamond$ then $E \vdash \diamond$.
2. If $E \triangleright T$ then $E \vdash T$.
3. If $E \triangleright S <: T$ and $E \vdash S$ then $E \vdash S <: T$.
4. If $E \vdash e \rightarrow T$ then $E \vdash e : T$.
5. If $E \vdash e \leftarrow T$ then $E \vdash e : T$.

Proof

By simultaneous induction over the derivations. For space reasons we give just the more interesting cases.

Part (4): (Synth Dot)

$$\frac{E \vdash e \rightarrow T \quad \text{norm}(T) = D \quad D.\ell \rightsquigarrow U}{E \vdash e.\ell \rightarrow [e.\ell : U]}$$

By induction hypothesis we have that $E \vdash e : T$. From Lemma 47 we have that $E \vdash T <: D$, and from Lemma 48 we have that $E \vdash D <: \{\ell : U\}$. By transitivity (Lemma 25) and the derived rule (Exp Subsum) we can conclude that $E \vdash e : \{\ell : U\}$. From rule (Exp Dot) we deduce that $E \vdash e.\ell : U$, and by the derived rule (Exp Singleton) we have $E \vdash e.\ell : [e.\ell : U]$ as required.

Part (4): (Synth Add)

$$\frac{E \vdash e_1 \rightarrow T_1 \quad E \vdash e_2 \rightarrow T_2 \quad \text{norm}(T_2) = D_2 \quad D_2.\text{Items} \rightsquigarrow U_2}{E \vdash e_1 :: e_2 \rightarrow ([e_1 : T_1] \mid U_2)^*}$$

By induction hypothesis we have both $E \vdash e_1 : T_1$ and $E \vdash e_2 : T_2$. It is simple to show that $E \vdash [e_1 : T_1] <: ([e_1 : T_1] \mid U_2)$ and so by rule (Exp Singular Subsum)

we can derive $E \vdash e_1 : ([e_1 : T_1] \mid U_2)$. From Lemma 47 we have that $E \vdash T_2 <: D_2$ and from Lemma 49 we have that $E \vdash D_2 <: U_2^*$. By transitivity (Lemma 25) and the derived rule (Exp Subsum) we can conclude that $E \vdash e_2 : U_2^*$. It is simple to show that $E \vdash U_2^* <: ([e_1 : T_1] \mid U_2)^*$ and by the derived rule (Exp Subsum) we can conclude that $E \vdash e_2 : ([e_1 : T_1] \mid U_2)^*$. From rule (Exp Add) we deduce $E \vdash e_1 :: e_2 : ([e_1 : T_1] \mid U_2)^*$ as required.

Part (5): (Swap)

$$\frac{E \vdash e \rightarrow T \quad E \triangleright [e : T] <: T'}{E \vdash e \leftarrow T'}$$

By (simultaneous) induction hypothesis we have that $E \vdash e : T$ and $E \vdash [e : T] <: T'$. By rule (Exp Singular Subsum) we have $E \vdash e : T'$ as required. \square

Appendix C.6 Exploiting SMT models correct

In this section we show that the operational checks we use to validate the models produced by the SMT solver are correct (Lemmas 51 and 52), and that the **elementof** construct does indeed return a value of the requested type or **null** (Lemma 54).

Lemma 51

If the three checks from Section 7.1 succeed then $E \not\vdash T <: T'$.

Proof

By inverting rule (Subtype), it suffices to show that $\models (\mathbf{F}[[E]] \wedge \mathbf{F}[[T]](x)) \Rightarrow \mathbf{F}[[T']](x)$. Since our intended model is not inconsistent, it suffices to show that

$$\models \exists x, y_1, \dots, y_n, \mathbf{F}[[U_1]](y_1) \wedge \dots \wedge \mathbf{F}[[U_n]](y_n) \wedge \mathbf{F}[[T]](x) \wedge \neg \mathbf{F}[[T']](x).$$

From conditions (1) and (2) by Lemma 2 it follows that $\models \mathbf{R}[[y_i \sigma \text{ in } U_i \sigma]] = \mathbf{true}$ for all $i \in 1..n$. After unfolding the definitions, this implies that $\models \mathbf{F}[[U_i \sigma]](y_i \sigma)$ for all $i \in 1..n$. In a similar way, from conditions (1) and (3) by Lemma 2 we have that $\models \mathbf{F}[[T \ \&! \ T'] \sigma](x \sigma)$, or equivalently that $\models \mathbf{F}[[T \sigma]](x \sigma) \wedge \neg \mathbf{F}[[T' \sigma]](x \sigma)$. Instantiating the existential variables with the values given by σ completes the proof. \square

Lemma 52

If the three checks in Section 7.1 succeed for $T' = \text{Empty}$ then $\emptyset \vdash x \sigma : T \sigma$ and $\emptyset \vdash y \sigma : U \sigma$ for all $(y : U) \in E$.

Proof

Since $x \sigma$ and $y \sigma$ for all $y \in \text{dom}(E)$ are values, by Lemma 31 and (Exp Singular Subsum) it suffices to show that $\emptyset \vdash [x \sigma] <: T \sigma$ and $\emptyset \vdash [y \sigma] <: U \sigma$ for all $y : U \in E$. By (Subtype) it suffices to show that $\models \mathbf{F}[[T \sigma]](x \sigma)$ and $\models \mathbf{F}[[U \sigma]](y \sigma)$ for all $y : U \in E$. These follow from the corresponding checks by Lemma 2 and basic reasoning in FOL. \square

Lemma 53

If $E \vdash T$ then the expression $v \text{ in } T$ is algorithmically pure.

Lemma 54

If **elementof** $T \rightarrow v$ and $\emptyset \vdash T$ then $\emptyset \vdash v : T \mid [\mathbf{null}]$.

Proof

By Lemma 53 we have that v in T is pure. By the reduction relation for **elementof** T we have that either $v = \mathbf{null}$, in which case the conclusion is immediate, or we know that v in $T \rightarrow^* \mathbf{true}$, which allows us to apply Lemma 52 for $E = \emptyset$ and obtain $\emptyset \vdash v : T$. \square

Acknowledgments

We thank Nikolaj Bjørner for his invaluable help in using Z3. James Margetson helped with F[#] programming issues. Paul Anderson, Ioannis Baltopoulos, Johannes Borgström, Nate Foster, Tim Harris, Ranjit Jhala, and Thorsten Tarrach commented on drafts. Discussions with Martín Abadi, Cliff Jones, Christoph Koch, Benjamin Pierce, and Dominique Unruh were useful, as were the comments of anonymous reviewers. Cătălin Hrițcu was supported by a fellowship from Microsoft Research and the International Max Planck Research School for Computer Science.

References

- Abiteboul, S., Buneman, P. & Suciu, D. (2000) *Data on the Web*. San Francisco, CA: Morgan Kaufmann.
- Aiken, A. & Wimmers, E. (1993) Type inclusion constraints and type inference. In the *Proceedings of ICFP 03, the Eighth ACM SIGPLAN International Conference on Functional Programming*, ACM, New York, USA.
- Aiken, A., Wimmers, E. L. & Lakshman, T. K. (1994) Soft typing with conditional types. In the *Proceedings of the 21st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, ACM, New York, USA.
- Aspinall, D. 1994 Subtyping with singleton types. In the *Proceedings of CSL*, LNCS vol. 933. Heidelberg, Germany: Springer.
- Aspinall, D. & Hofmann, M. (2005) Dependent types. In *Advanced Topics in Types and Programming Languages*, chapter 2. Cambridge, MA: MIT Press.
- Backes, M., Hrițcu, C. & Tarrach, T. (2011) Automatically verifying typing constraints for a data processing language. In the *Proceedings of CPP, the 11th Generative Approaches to Second Language Acquisition Conference (GASLA 2011)*. Somerville, MA: CPP.
- Barnett, M., Chang, B.-Y. E., DeLine, R., Jacobs, B. & Leino, K. R. M. 2005 Boogie: A modular reusable verifier for object-oriented programs. In the *Proceedings of FMC0*, Amsterdam, The Netherlands, LNCS vol. 4111.
- Barrett, C., Deters, M., Oliveras, A. & Stump, A. (2008) Design and results of the 3rd Annual SMT Competition (SMT-COMP 2007). *Int. J. Artif. Intell. Tools.* **17**(4), 569–606.
- Barrett, C. & Tinelli, C. (2007) CVC3. In the *Proceedings of CAV*, Berlin, Germany, LNCS vol. 4590.
- Bengtson, J., Bhargavan, K., Fournet, C., Gordon, A. D. & Maffei, S. (2008) Refinement types for secure implementations. In the *Proceedings of CSF*, Pittsburgh, PA, USA.
- Benzaken, V., Castagna, G. & Frisch, A. (2003) CDuce: An XML-friendly general purpose language. In the *Proceedings of the Eighth ACM SIGPLAN International Conference on Functional Programming (ICFP)*, Uppsala, Sweden.
- Bierman, G. M., Gordon, A. D., Hrițcu, C. & Langworthy, D. (2010a) *Semantic Subtyping with an SMT Solver*. Technical Report MSR-TR-2010-99, Microsoft Research, Redmond, WA.

- Bierman, G. M., Gordon, A. D., Hrițcu, C. & Langworthy, D. (2010b) Semantic subtyping with an SMT solver. In the *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, Baltimore, MD, USA.
- Bierman, G. M., Meijer, E. & Torgersen, M. (2007) Lost in translation: Formalizing proposed extensions to C#. In the *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming (OOPSLA)*, Baltimore, MD, USA.
- Böhme, S., Leino, K. R. M. & Wolff, B. (2008) HOL-Boogie – an interactive prover for the Boogie program-verifier. In the *Proceedings of TPHOLs*, Montreal, Canada, LNCS vol. 5170.
- Box, D. (2010) Update on SQL Server Modeling CTP (Repository/Modeling Services, “Quadrant” and “M”). Accessed September 22, 2010. Blog available at <http://blogs.msdn.com/b/modelcitizen>
- Boyer, R. S., Kaufmann, M. & Moore, J. S. (1995) The Boyer-Moore theorem prover and its interactive enhancement. *Comput. Math. Appl.* **29**(2), 27–62.
- Buneman, P., Naqvi, S., Tannen, V. & Wong, L. (1995) Principles of programming with complex objects and collection types. *Theor. Comput. Sci.* **149**(1), 3–48.
- Buneman, P. & Pierce, B. C. (1999) Union types for semi-structured data. In the *Proceedings of DBPL*, Kinloch Rannoch, Scotland, UK, LNCS vol. 1949.
- Burstall, R. M., MacQueen, D. B. & Sannella, D. (1980) HOPE: An experimental applicative language. In the *Proceedings of LISP Conference*, Stanford, CA, USA.
- Calcagno, C., Cardelli, L. & Gordon, A. D. (2005) Deciding validity in a spatial logic for trees. *J. Funct. Program.* **15**, 543–572.
- Cartwright, R. & Fagan, M. (1991) Soft typing. In the *Proceedings of PLDI*, Toronto, Ontario, Canada.
- Castagna, G. (2005) Patterns and types for querying XML documents. In the *Proceedings of DBPL*, Trondheim, Norway, LNCS vol. 3774.
- Castagna, G. & Chen, G. (2001) Dependent types with subtyping and late-bound overloading. *Inf. Comput.* **168**(1), 1–67.
- Cohen, S. (2006) User-defined aggregate functions: Bridging theory and practice. In the *Proceedings of SIGMOD*, Chicago, IL, USA.
- Crockford, D. (2006) *The Application/Json Media Type for JavaScript Object Notation (JSON)*. RFC 4627. Reston, VA: The Internet Society.
- Damm, F. (1994) Subtyping with union types, intersection types and recursive types. In the *Proceedings of TACS*, Sendai, Japan.
- de Moura, L. M. & Bjørner, N. (2007) Efficient E-matching for SMT solvers. In the *Proceedings of CADE-21*, Bremen, Germany, LNCS vol. 4603. Heidelberg, Germany: Springer, pp. 183–198.
- de Moura, L. M. & Bjørner, N. (2008) Z3: An efficient SMT solver. In the *Proceedings of TACAS*, Budapest, Hungary.
- de Moura, L. M. & Bjørner, N. (2009) Generalized, efficient array decision procedures. In *Proceedings of FMCAD*, Austin, TX, USA.
- Detlefs, D., Nelson, G. & Saxe, J. B. (2005) Simplify: A theorem prover for program checking. *J. ACM.* **52**(3), 365–473.
- Dunfield, J. (Aug. 2007) *A Unified System of Type Refinements*. PhD. thesis, CMU-CS-07-129, Carnegie Mellon University, Pittsburgh, PA.
- Dunfield, J. & Pfenning, F. (2004) Tridirectional typechecking. In the *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL*, Venice, Italy.
- Dutertre, B. & de Moura, L. M.. The YICES SMT solver. Accessed February 27, 2012. Available at: <http://yices.csl.sri.com/tool-paper.pdf>, 2006.
- Findler, R. & Felleisen, M. (2002) Contracts for higher-order functions. In the *Proceedings of the Seventh ACM SIGPLAN International Conference on Functional Programming (ICFP '02)*, Pittsburgh, PA, USA.

- Fisher, K., Mandelbaum, Y. & Walker, D. (2006) The next 700 data description languages. In the *Proceedings of the Symposium on Principles of Programming Languages (POPL)*, Charleston, SC.
- Flanagan, C. (2006) Hybrid-type checking. In the *Proceedings of the Symposium on Principles of Programming Languages (POPL)*, Charleston, SC.
- Freeman, T. & Pfenning, F. (1991) Refinement types for ML. In the *Proceedings of the ACM SIGPLAN'91 Conference on Programming Language Design and Implementation (PLDI)*, Toronto, Ontario, Canada.
- Frisch, A., Castagna, G. & Benzaken, V. (2008) Semantic subtyping: Dealing set-theoretically with function, union, intersection, and negation types. *J. ACM* **55**(4), 19:1–19:64.
- Genevès, P., Layaida, N. & Schmitt, A. (2007) Efficient static analysis of XML paths and types. In the *Proceedings of the ACM SIGPLAN 2007 Conference on Programming Language Design and Implementation (PLDI)*, San Diego, CA, USA.
- Giesl, J. (1997) Termination of nested and mutually recursive algorithms. *J. Autom. Reasoning* **19**, 1–29.
- Gordon, A. D. & Jeffrey, A. (2002) Typing one-to-one and one-to-many correspondences in security protocols. In the *Proceedings of ISSS*, Kyoto, Japan.
- Greenberg, M., Pierce, B. & Weirich, S. (2010) Contracts made manifest. In the *Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, Madrid, Spain.
- Greve, D. A., Kaufmann, M., Manolios, P., Moore, J. S., Ray, S., Ruiz-Reina, J.-L., Sumners, R., Vroon, D. & Wilding, M. (2008) Efficient execution in an automated reasoning environment. *J. Funct. Program.* **18**(1), 15–46.
- Hosoya, H. & Pierce, B. (2003) XDuce: A statically typed XML processing language. *ACM Trans. Internet Technol.* **3**(2), 117–148.
- Hosoya, H., Vouillon, J. & Pierce, B. (2000) Regular expression types for XML. In the *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP)*, Montreal, Canada.
- Jaffar, J. & Maher, M. J. (1994) Constraint logic programming: A survey. *J. Log. Algebr. Program.* **19**(20), 503–581.
- Jhala, R., Majumdar, R. & Rybalchenko, A. (2011) HMC: Verifying functional programs using abstract interpreters. In the *Proceedings of CAV*, Snowbird, UT, USA. DBLP, pp. 470–485.
- Jhala, R., Majumdar, R. & Xu, R.-G. (2007) State of the union: Type inference via Craig interpolation. In the *Proceedings of TACAS*, Braga, Portugal.
- Jones, C. (1986) *Systematic Software Development Using VDM*. Upper Saddle River, NJ: Prentice-Hall.
- Kawaguchi, M., Rondon, P. M. & Jhala, R. (2009) Type-based data structure verification. In the *Proceedings of the 2009 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Dublin, Ireland. New York: ACM, pp. 304–315.
- King, J. C. (1976) Symbolic execution and program testing. *Commun. ACM.* **19**, 385–394.
- Knowles, K., Tomb, A., Gronski, J., Freund, S. & Flanagan, C. (2007) SAGE: *Unified Hybrid Checking for First-Class Types, General Refinement Types and Dynamic*. Technical Report, UCSC, Santa Cruz, CA.
- Knowles, K. W. & Flanagan, C. (2010) Hybrid type checking. *ACM TOPLAS* **32**(2), 6:1–6:34.
- Komondoor, R., Ramalingam, G., Chandra, S. & Field, J. (2005) Dependent types for program understanding. In the *Proceedings of the Eleventh International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, Edinburgh, UK.
- Kopylov, A. (2003) Dependent intersection: A new way of defining records in type theory. In the *Proceedings of the 18th IEEE Symposium on Logic in Computer Science (LICS)*, Ottawa, Canada.
- Kuncak, V., Mayer, M., Piskac, R. & Suter, P. (2010) Complete functional synthesis. In the *Proceedings of the 2010 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*. New York: ACM, pp. 316–329.

- Leino, K. R. M. & Monahan, R. (2009) Reasoning about comprehensions with first-order SMT solvers. In the *Proceedings of the ACM Symposium on Applied Computing (SAC)*, Honolulu, HI, USA.
- Lerner, B., Flower, M., Grossman, D. & Chambers, C. (2007) Searching for type-error messages. In the *Proceedings of PLDI*, San Diego, CA, USA.
- Lovas, W. & Pfenning, F. (2007) A bidirectional refinement type system for LF. In the *Proceedings of LFMTP*, Bremen, Germany.
- McCarthy, J. (1962) Towards a mathematical science of computation. In the *Proceedings of IFIP Congress*, Munich, Germany.
- Meijer, E., Beckman, B. & Bierman, G. M. (2007) LINQ: Reconciling objects, relations and XML in the .NET framework. In the *Proceedings of SIGMOD*, Beijing, China.
- Meng, J. & Paulson, L. C. (2008) Translating higher-order problems to first-order clauses. *J. Autom. Reasoning* **40**(1), 35–60.
- Meyer, B. 1992 *Eiffel: The Language*. Upper Saddle River, NJ: Prentice Hall.
- Microsoft Corporation (Oct. 2009) *The Microsoft Code Name “M” Modeling Language Specification Version 0.5*. Preliminary implementation available as part of the SQL Server Modeling CTP (November 2009). Redmond, WA: Microsoft Corp.
- Nordström, B. & Petersson, K. (1983) Types and specifications. In the *Proceedings of IFIP*, Paris, France.
- Pierce, B. C. (1991) *Programming with Intersection Types, Union Types, and Polymorphism*. Technical Report CMU-CS-91-106, Carnegie Mellon University, Pittsburgh, PA.
- Pierce, B. C. (2002) *Types and Programming Languages*. Cambridge, MA: MIT Press.
- Pierce, B. C. & Turner, D. N. (2000) Local type inference. *ACM Trans. Program. Lang. Syst.* **22**(1), 1–44.
- Pratt, V. (1983) Five paradigm shifts in programming language design and their realization in Viron, a dataflow programming environment. In the *Proceedings of POPL*, Austin, TX, USA.
- Ranise, S. & Tinelli, C. (2006) *The SMT-LIB Standard: Version 1.2*. Technical Report, Department of Computer Science, The University of Iowa, Iowa.
- Reynolds, J. C. (1996) Design of the programming language Forsythe. In *Algol-Like Languages*, O’Hearn, P. W. & Tennant, R. D. (eds), chapter 8. Basel, Switzerland: Birkhäuser, pp. 173–234.
- Rondon, P., Kawaguchi, M. & Jhala, R. (2008) Liquid types. In the *Proceedings of PLDI*, Tucson, AZ, USA.
- Rushby, J., Owre, S. & Shankar, N. (1998) Subtypes for specifications: Predicate subtyping in PVS. *IEEE Trans. Softw. Eng.* **24**(9), 709–720.
- Saraswat, V., Nystrom, N., Palsberg, J. & Grothoff, C. (2008) Constrained types for object-oriented languages. In the *Proceedings of OOPSLA*, Nashville, TN, USA.
- Siméon, J. & Wadler, P. (2003) The essence of XML. In the *Proceedings of POPL*, New Orleans, USA.
- Sozeau, M. (2006) Subset coercions in Coq. In the *Proceedings of TYPES*, Nottingham, UK, LNCS vol. 4502.
- Swamy, N., Chen, J. & Chugh, R. (2010) Enforcing stateful authorization and information flow policies in Fine. In the *Proceedings of ESOP*, Paphos, Cyprus.
- Terauchi, T. (2010) Dependent types from counterexamples. In the *Proceedings of POPL*, Madrid, Spain.
- Tobin-Hochstadt, S. & Felleisen, M. (2008) The design and implementation of Typed Scheme. In the *Proceedings of the 35th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, San Francisco, CA, USA.
- Tobin-Hochstadt, S. & Felleisen, M. (2010) Logical types for untyped languages. In the *Proceedings of the 15th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, Baltimore, MD, USA.

- TypiCal Project 2009 *The Coq Proof Assistant*. Version 8.2. Accessed February 27, 2012. Available at: <http://coq.inria.fr>.
- Unno, H. & Kobayashi, N. (2009) Dependent type inference with interpolants. In the *Proceedings of the 11th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP)*, Coimbra, Portugal.
- Wright, A. K. & Cartwright, R. (1997) A practical soft type system for scheme. *ACM TOPLAS* **19**, 87–152.
- Wright, A. K. & Felleisen, M. (1994) A syntactic approach to type soundness. *Inf. Comput.* **115**(1), 38–94.
- Xi, H. & Pfenning, F. (1999) Dependent types in practical programming. In the *Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, San Antonio, TX, USA.