

Structural and Flow-Sensitive Types for Whiley

David J. Pearce and James Noble

School of Engineering and Computer Science
Victoria University of Wellington, NZ
djp@ecs.vuw.ac.nz

July 2011

Abstract

Modern statically typed languages require variables to be declared with a single static type, and that subtyping relationships between user-defined types be made explicit. This contrasts with dynamically typed languages, where variables are declared implicitly, can hold values of different types at different points and have no restrictions on flow (leading to ad-hoc and implicit subtyping).

We present the flow-sensitive and structural type system used in the Whiley language. This permits variables to be declared implicitly, have multiple types within a function, and be retyped after runtime type tests. Furthermore, subtyping between user-defined types is implicit, based purely on structure. The result is a statically-typed language which, for the most part, has the look and feel of a dynamic language. The typing algorithm operates in a fashion similar to dataflow analysis. Widening must be applied to ensure termination although, surprisingly, there is no loss of precision. We formalise Whiley’s type system and operational semantics, and give proofs of termination and soundness.

1 Introduction

Statically typed programming languages lead to programs which are more efficient and where errors are easier to detect ahead-of-time [1, 2]. Static typing forces some discipline on the programming process. For example, it ensures at least some documentation regarding acceptable function inputs is provided. In contrast, dynamically typed languages are more flexible in nature which helps reduce overheads and increase productivity [3, 4, 5, 6]. Indeed, recent times have seen a significant shift towards dynamically typed languages [7].

Numerous attempts have been made to bridge the gap between static and dynamic languages. Scala [8], C#3.0 [9], OCaml [10] and, most recently, Java 7 all employ local type inference (in some form) to reduce syntactic overhead. Techniques such as gradual typing [11, 12], soft typing [1, 13] and hybrid typing [14] enable a transitory position where some parts of a program are statically typed, and others are not. Alternatively, global type inference can be used (in some situations) to reconstruct types “after the fact” for programs written in dynamic languages [15, 16].

Such approaches focus on working around static type systems, rather than rethinking the way static typing is done. The prevailing view remains that each variable in a static type system must have exactly one declared type and that *nominal* — rather than *structural* — typing is preferred. And yet, these choices go against the notion of a dynamic language, where variables can be assigned arbitrarily and can flow unimpeded by cumbersome, rigid type hierarchies.

We present the type system used in Whiley, a statically-typed language targeting the JVM. Whiley takes a novel approach to typing, which we refer to as *flow-sensitive typing*. This comes close to giving the flexibility of a dynamic language, with the guarantees of a statically typed language. More specifically, flow-sensitive typing offers improved error handling, reduced syntactic overhead and greater opportunity for code reuse. The technique is adopted from flow-sensitive program analysis (e.g. [17, 18, 19]), and allows variables to have different types at different points and be declared implicitly, based on the type of assigned expressions. Furthermore, unlike the majority of statically typed languages, Whiley employs structural — rather than nominal — typing. This frees programmers from the burden of developing static type hierarchies up front, allowing them to expose subtyping relationships retroactively.

Finally, an open source implementation of Whiley is freely available from <http://whiley.org> and details Whiley’s JVM implementation can be found here [20].

1.1 Flow-Sensitive Types

The following demonstrates Whiley’s flow-sensitive types:

```
define Circle as {int x, int y, int r}
define Rect as {int x, int y, int w, int h}
define Shape as Circle | Rect

real area(Shape s):
  if s is Circle:
    return PI * s.r * s.r
  else:
    return s.w * s.h
```

A `Shape` is either a `Rect` or a `Circle` (which are both record types). The type test “`s is Circle`” (similar to an `instanceof` test in Java) determines whether `s` is a `Circle` or not. Unlike Java, Whiley automatically retypes `s` to have type `Circle` (resp. `Rect`) on the true (resp. false) branches of the `if` statement. There is no need to explicitly cast `s` to the appropriate `Shape` before accessing its fields.

Statically typed languages typically require variables be explicitly declared. Compared with dynamically typed languages, this is an extra burden for the programmer, particularly when a variable’s type can be inferred from assigned expression(s). In Whiley, local variables are never explicitly declared, rather they are declared by assignment:

```
int average([int] items):
  v = 0
  for i in items:
    v = v + items[i]
  return v / |items|
```

Here, `items` is a list of `ints`, whilst `|items|` returns its length. Variable `v` accumulates the sum of all elements in the list, and is declared by the assignment “`v = 0`”. Since `0` has type `int`, `v` has type `int` after the assignment.

1.2 Structural Subtyping

Statically typed languages, such as Java, employ nominal typing for recursive data types. This results in rigid hierarchies which are often difficult to extend [21]. In contrast, Whiley employs *structural subtyping* of records [22] to give greater flexibility. For example, the following defines a `Border` record:

```
define Border as {int x,int y,int w,int h}
```

Any instance of `Border` has identical structure to an instance of `Rect`. Wherever a `Border` is required, a `Rect` can be provided and vice-versa — even if the `Border` definition was written long after the `Rect`, and even though no explicit connection is made between `Rect` and `Border`.

The focus on structural, rather than nominal, types in Whiley is also evident in the way instances are created:

```
bool contains(int x, int y, Border b):
  ....

bool example(int x, int y):
  b = {x: 1, y: 2, w: 10, h: 3}
  return contains(x,y,b)
```

Here, function `example()` creates a record instance with fields `x`, `y`, `w` and `h`, and assigns it to variable `b`. Despite not being associated with a name, such as `Border` or `Rect`, it can be freely passed into functions expecting such types, since they have identical *structure*.

2 Language Overview

Figure 1 provides a simple implementation of expressions, along with code for evaluating them. The types `Expr` and `Value` are algebraic data types, with the latter defining the set of allowed values. Type `Op` is an enumeration, whilst `BinOp` and `ListAccess` are records which form part of `Expr`. Parameter `env` is a map from variables to `Values`. Finally, `null` is used as an error condition to indicate a “stuck” state (i.e. the evaluation cannot proceed).

The code in Figure 1 makes extensive use of runtime type tests to distinguish different expression forms (e.g. “`e is int`”). These work in a similar fashion to Java’s `instanceof` operator, with one important

```

1 define Var as string
2 define Op as { ADD, SUB, MUL, DIV }
3 define BinOp as { Op op, Expr lhs, Expr rhs }
4 define ListAccess as { Expr lhs, Expr rhs }
5
6 define Value as int | [Value] | null
7
8 define Expr as int | Var | BinOp | [Expr] | ListAccess
9
10 Value eval(Expr e, {Var→Value} env):
11   if e is int:
12     return e
13   else if e is Var && e in env:
14     // look up variable's value
15     return env[e]
16   else if e is BinOp:
17     // evaluate left and right expressions
18     lhs = eval(e.lhs, env)
19     rhs = eval(e.rhs, env)
20     // sanity check
21     if !(lhs is int && rhs is int):
22       return null // stuck
23     // evaluate result
24     switch e.op:
25       case ADD:
26         return lhs + rhs
27       case SUB:
28         return lhs - rhs
29       case MUL:
30         return lhs * rhs
31       case DIV:
32         if rhs != 0:
33           return lhs / rhs
34   else if e is ListAccess:
35     // evaluate src and index expressions
36     src = eval(e.lhs, env)
37     index = eval(e.rhs, env)
38     // sanity check
39     if src is [Value] && index is int
40         && index >= 0 && index < |src|:
41       return src[index]
42   else if e is [Expr]:
43     lv = []
44     // evaluate items in list constructor
45     for i in e:
46       v = eval(i, env)
47       if v == null:
48         return v
49       else:
50         lv = lv + [v]
51   return lv
52 // some kind of error occurred, so propagate upwards
53 return null

```

Figure 1: Whiley code for a simple expression tree and evaluation function. This makes extensive use of type tests, both for distinguishing expressions and error handling. Flow-sensitive typing greatly simplifies the code, which would otherwise require numerous unnecessary casts.

difference: they operate in a flow-sensitive fashion and automatically *retype* variables after the test. As an example, consider the type test “`e is int`” on Line 11. On the true branch, variable `e` is automatically retyped to have type `int`. Likewise, on the false branch, `e` is now known *not* to have type `int` (and any attempt to retest this yields a compile-time error).

Figure 1 also employs runtime type tests to identify and propagate errors. For example, having evaluated the left- and right-hand sides of a `BinOp`, we check on Line 21 that both are `int` values (i.e. not list values or `null`). After the check, Whiley’s flow-sensitive type system automatically retypes both `lhs` and `rhs` to `int`. For `ListAccess` expressions, we check on Line 39 that `src` is a list value, and that `index` is an `int`. The latter is achieved with “`index is int`”. As `src` is retyped within the condition itself, the subsequent use of `|src|` on Line 40 is type safe.

Implementing our expression language in a statically-typed language, such as Java, would require code that was more cumbersome, and more verbose than that of Figure 1. One reason for this is that, in languages like Java, variables must be *explicitly* retyped after `instanceof` tests. That is, we must insert casts to update the types of tested variables and, since variables can have only one type in Java, introduce temporary variables to hold these new types. For example, after a test “`e instanceof BinOp`” we must introduce a new variable, say `r`, with type `BinOp` and assign `e` to `r` using an appropriate cast. A Java implementation would also (most likely) break up the test on Line 39, since it would otherwise need two identical casts (one inside the condition for `|src|`, and one on the true branch for `src[index]`).

In an object-oriented language, such as Java, a direct conversion of Figure 1 might not be optimal. Instead, the *visitor pattern* [23] can be used to distinguish different expression forms. Using the visitor pattern reduces the amount of explicit retyping required. This is because the different expression forms are explicitly given as parameters to the visitor methods. However, using the visitor pattern is a heavyweight solution which is not suitable in all situations. In particular, it would not eliminate all forms of explicit retyping from Figure 1. In this case, explicit variable retyping will still be required to properly handle the different values returned from `eval()`. For example, to check `BinOps` are evaluated on `int` operands (Line 21), and that `src` gives a list and `index` an `int` (Line 39).

2.1 Value Semantics

In Whiley, all compound structures (e.g. lists, sets, and records) have *value semantics*. This means they are passed and returned by-value (as in Pascal, MATLAB or most functional languages). But, unlike functional languages (and like Pascal), values of compound types can be updated in place.

Value semantics implies that updates to a variable only affects that variable, and that information can only flow out of a function through its return value. Whiley has no general, mutable heap comparable to those found in object-oriented languages. Consider:

```
int f([int] xs):
  ys = xs
  xs[0] = 1
  ...
```

The semantics of Whiley dictate that, having assigned `xs` to `ys` as above, the subsequent update to `xs` does not affect `ys`. Arguments are also passed by value, hence `xs` is updated inside `f()` and this does not affect `f`’s caller. That is, `xs` is not a *reference* to a list of `int`; rather, it *is* a list of `ints` and assignments to it do not affect state visible outside of `f()`.

Whilst this approach may seem inefficient, a variety of techniques exist (e.g. reference counting) to ensure efficiency (see e.g. [24, 25, 26]). Indeed, the underlying implementation does pass compound structures by reference and copies them only when absolutely necessary.

2.1.1 Structural Updates.

In a dynamic language, lists and records can be updated at will. However, static type systems normally require updates to respect the element or field type in question. For example, assigning a `float` to an element of an `int` array is not permitted in Java. To work around this, programmers typically either clone the structure in question, or “break” the type system using casting (or similar).

Updates to list elements and record fields are always permitted in Whiley. For example:

```
define Point as {int x, int y}

[int|Point] insert(int i, Point p, [int] xs):
  xs[i] = p
  return p
```

The type of `xs` is updated to `[int|Point]` after `xs` is assigned as it now contains a `Point`.

The ability to update the types of records and lists is possible because compound structures (e.g. lists, sets, records, etc) have value semantics in Whiley. Such updates would be unsafe if aliasing were permitted as, for example, callers of `insert()` might hold aliases to `xs` and would expect its elements to be `ints`.

2.2 Structural Types

Statically typed languages, such as Java, employ nominal typing for recursive data types. This results in rigid hierarchies which are often difficult to extend [21]. In contrast, Whiley employs *structural subtyping* of records [22] to give greater flexibility.

Suppose we wish to extend our expressions from Figure 1 with assignment statements. A common issue arises as the left-hand side of an assignment is a restricted form of expression, often called an *lval*. In a language like Java, we can capture this nicely using interfaces:

```
interface Expr { ... }
interface LVal { ... }
class ListAccess implements Expr, LVal { ... }
class Var implements Expr, LVal { ... }
class Int implements Expr { ... }
```

However, suppose the code for expressions was part of an existing library, and we are trying to add statements after the fact. In a language like Java, this presents a problem as we cannot retroactively insert the necessary `LVal` interface to `Var` and `ListAccess`.

In Whiley, adding the notion of an `LVal` is easy to do retroactively because of structural subtyping:

```
define LVal as Var | ListAccess
define Assign as {LVal lhs, Expr rhs}
```

```
Expr parseExpression():
  ...
```

```
null|Assign parseAssign():
  le = parseExpression()
  match(":=")
  re = parseExpression()
  if le is LVal:
    return {lhs: le, rhs: re}
  else:
    return null //syntax error
```

Here, `LVal` is implicitly a subtype of `Expr` — i.e. there is no need for an explicit declaration of this, as would be required in Java. That is, they can be defined entirely separately from each other (e.g. in different files, packages or entirely separate programs) — and yet, `LVal` remains a subtype of `Expr`.

2.2.1 Structural Subtyping.

Whiley permits subtyping between recursive structural types.

```
define Link as {int data, LinkedList next}
define LinkedList as null | Link
define OrderedList as null | {
  int data, int order, OrderedList next
}

int sum(LinkedList l):
  if l is null:
    return 0
  else:
    return l.data + sum(l.next)
```

Here, we have defined a standard linked list and a specialised “ordered” list where `order < next.order` for each node (see e.g. [27]). Whiley type checks this function by showing that `OrderedList` is a structural subtype of `LinkedList` — despite this relationship not being identified explicitly in the program. Type checking in the presence of recursive structural types is a well-known and challenging problem [28, 29, 30] which is further compounded in Whiley by the presence of flow-sensitive reasoning.

2.3 Flow-Sensitive Types

We now examine Whiley’s flow-sensitive type system in more detail, and identify several ways in which it leads to improved code quality.

Error Handling. Nullable references have proved a significant source of error in e.g. Java [31]. The issue is that, in such languages, one can treat *nullable* references as though they are *non-null* references [32]. Many solutions have been proposed which distinguish these two forms using static type systems (e.g. [33, 34, 35, 36, 37, 38, 39, 40]).

Whiley’s flow-sensitive type system lends itself naturally to handling this problem because it supports *union types* (e.g. [41, 42]). These allow variables to hold values from different types, rather than just one type. For example:

```

null|int indexOf(string str, char c):
    ...

[string] split(string str, char c):
    idx = indexOf(str,c)
    // idx has type null|int
    if idx is int:
        // idx now has type int
        below = str[0..idx]
        above = str[idx..]
        return [below,above]
    else:
        // idx now has type null
        return [str] // no occurrence
```

Here, `indexOf()` returns the first index of a character in the string, or `null` if there is none. The type `null|int` is a union type, meaning it is either an `int` or `null`.

In the above example, Whiley’s flow-sensitive type system seamlessly ensures that `null` is never dereferenced. This is because the type `null|int` cannot be treated as an `int`. Instead, one must first check it *is* an `int` using a type test, such as “`idx is int`”. Whiley automatically retypes `idx` to `int` when this is known to be true, thereby avoiding any awkward and unnecessary syntax (e.g. a cast as required in e.g. [43, 38]).

Code Reuse. Whiley’s flow-sensitive type system can expose greater opportunities for code reuse:

```

1{string} usedVariables(Expr e):
2    if e is Var:
3        return {e}
4    else if e is BinOp || e is ListAccess:
5        l = useVariables(e.lhs)
6        r = useVariables(e.rhs)
7        return l + r // set union
8    else if e is [Expr]:
9        ...
10   else:
11       return {}
```

On Line 5, variable `e` has type `BinOp|ListAccess`. The use of `e.lhs` at this point is type safe, since we can perform operations common to all types of a union and, in particular, unions of records expose common fields (similar to a *common initial sequence* for unions of structs in C [44, §6.3.2.3]).

In languages like Java, exploiting code reuse in this way requires careful planning, as common types must be explicitly related in the class hierarchy. In contrast, Whiley’s flow-sensitive type system lets us exploit opportunities for code reuse in an ad-hoc fashion, as and when they occur.

3 Type System

We formalise the Whyley language using a core calculus which, in the spirit of Featherweight Java [45], we call *Featherweight Whyley (FW)*. The following gives a syntactic definition of types in FW, where overbar (e.g. \bar{T}) indicates a list of items numbered consecutively (e.g. T_1, \dots, T_n):

$$T ::= \text{any} \mid \text{void} \mid \text{null} \mid \text{int} \mid \bar{T} \rightarrow T \mid [T] \mid \{\bar{T} \bar{n}\} \mid T_1 \vee T_2 \mid \mu X.T \mid X$$

Here, **void** and **any** represent \perp and \top ; lists are given by $[T]$; and, $\{\bar{T} \bar{n}\}$ represents records with one or more fields. The union $T_1 \vee T_2$ is a type whose values are in T_1 or T_2 . Types are additionally restricted to being *contractive* [29]. This prohibits types of the form $\mu X.X$ and $\mu X.(X \vee \dots)$ and is necessary to ensure a strong connection with regular trees [32].

FW follows the *equi-recursive* approach to dealing with recursive types [32]. That is, we do not distinguish between recursive types and their unfoldings. This simplifies FW's formalism, as one does not need an explicit *unfold* operator. Any implementation of FW, however, must address the fact that two equivalent types can have distinct representations on the machine (analogous to the graph isomorphism problem). We return to discuss this in §3.4.

3.1 Semantic Interpretation

To better understand the meaning of types in FW, it is helpful to give a *semantic interpretation* (following e.g. [46, 47, 48]). The aim is to give a set-theoretic model where subtype corresponds to subset. Whilst our model ties very closely with the semantics of FW, it is simplified to avoid a circular definition between types and semantics [48]. This is often referred to as “bootstrapping” the subtype relation, and affects only function values — which we model as sets of parameter and return values.

The following defines the language of values in our model:

$$V ::= \text{null} \mid i \mid [\bar{V}] \mid \{\bar{n} : \bar{V}\} \mid (\bar{V}_1, \dots, \bar{V}_n) \rightarrow \bar{V}$$

Here, $i \in \mathcal{I}$ represents integer values, whilst $(\bar{V}_1, \dots, \bar{V}_n) \rightarrow \bar{V}$ is our model of functions used to bootstrap the subtyping relation (note the abuse of notation here where \bar{V} indicates a set not a list — but, this should always be clear in context). Each function value consists of sets of parameter and return values which identify those values it can accept in each parameter position, and which it may return. There is no explicit connection between individual parameter values and the corresponding return value, as this level of detail is unnecessary.

Definition 1 (Type Acceptance) A type T accepts a value V , denoted by $T \models V$, defined as follows:

$$\begin{array}{ll} \text{any} \models V & \\ \text{null} \models \text{null} & \\ \text{int} \models i & \text{if } i \in \mathcal{I} \\ [T] \models [\bar{V}] & \text{if } T \models \bar{V} \\ \{T_1 \bar{n}_1, \dots, T_n \bar{n}_n\} \models \{\bar{n}_1 : V_1, \dots, \bar{n}_m : V_m\} & \text{if } n < m, T \models \bar{V} \\ T_1 \vee T_2 \models V & \text{if } T_i \models V, i \in \{1, 2\} \\ \bar{T} \rightarrow T \models (\bar{V}_1, \dots, \bar{V}_n) \rightarrow \bar{V} & \text{if } T \models \bar{V}, \bar{V}_1 \models T, \dots, \bar{V}_n \models T \end{array}$$

Here, $\bar{V} \models T$ denotes $\neg \exists V. [T \models V \wedge V \notin \bar{V}]$. In other words, every value accepted by T is in the set \bar{V} .

Definition 2 (Type Subset) We take $T_1 \models T_2$ to denote $\forall V. [T_1 \models V \implies T_2 \models V]$. In other words, that the set of values accepted by T_1 is a subset of those accepted by T_2 .

We are essentially viewing types as finite tree automata which accept conforming trees (i.e values). The rule for accepting function values may seem somewhat strange. However, consider this example:

```

int f (any x) :
  return 1

int g (int (int) fp) :
  return fp(123)

```

Here, fp is a variable of type $\text{int} \rightarrow \text{int}$, whilst function $f()$ has type $\text{any} \rightarrow \text{int}$. Intuitively, $g(\&f)$ should be a syntactically correct expression and, hence, $\text{any} \rightarrow \text{int}$ must subtype $\text{int} \rightarrow \text{int}$ (as is normal). Essentially, the type $\text{int} \rightarrow \text{int}$ will accept any function value which is prepared to accept any **int** value.

Subtyping:		
$\frac{}{T \leq T \mid \mathcal{C}}$	$\frac{\{T_1 \leq T_2\} \subseteq \mathcal{C}}{T_1 \leq T_2 \mid \mathcal{C}}$	(S-REFLEX, S-INDUCT)
$\frac{}{\text{void} \leq T \mid \mathcal{C}}$	$\frac{}{T \leq \text{any} \mid \mathcal{C}}$	(S-VOID, S-ANY)
$\frac{T_1 \leq T_2 \mid \mathcal{C} \cup \{[T_1] \leq [T_2]\}}{[T_1] \leq [T_2] \mid \mathcal{C}}$		(S-LIST)
$\begin{array}{l} n < m \\ T_1 = \{T_3 f_3, \dots, T_m f_m\} \\ T_2 = \{T'_3 f_3, \dots, T'_n f_n\} \\ T_3 \leq T'_3 \mid \mathcal{C} \cup \{T_1 \leq T_2\} \\ \dots \\ T_n \leq T'_n \mid \mathcal{C} \cup \{T_1 \leq T_2\} \end{array}$		(S-REC)
$\frac{\frac{\overline{T} \geq \overline{S} \mid \mathcal{C} \cup \{\overline{T} \rightarrow T \leq \overline{S} \rightarrow S\}}{T \leq S \mid \mathcal{C} \cup \{\overline{T} \rightarrow T \leq \overline{S} \rightarrow S\}}}{\overline{T} \rightarrow T \leq \overline{S} \rightarrow S \mid \mathcal{C}}$		(S-FUN)
$\frac{i \in \{2, 3\}}{T_1 \leq T_i \mid \mathcal{C} \cup \{T_1 \leq T_2 \vee T_3\}}{T_1 \leq T_2 \vee T_3 \mid \mathcal{C}}$		(S-UNION1)
$\frac{T_2 \leq T_1 \mid \mathcal{C} \cup \{T_2 \vee T_3 \leq T_1\}}{T_3 \leq T_1 \mid \mathcal{C} \cup \{T_2 \vee T_3 \leq T_1\}}{T_2 \vee T_3 \leq T_1 \mid \mathcal{C}}$		(S-UNION2)
$\frac{T = \{T_1 f_1 \dots, T_i \vee T'_i f_i \dots, T_n f_n\}}{\frac{\{T_1 f_1 \dots, T_i f_i \dots, T_n f_n\} \leq S \mid \mathcal{C} \cup \{T \leq S\}}{\{T_1 f_1 \dots, T'_i f_i \dots, T_n f_n\} \leq S \mid \mathcal{C} \cup \{T \leq S\}}}{T \leq S \mid \mathcal{C}}$		(S-UNION3)

Figure 2: Subtype rules for Featherweight Whieley. There is no need for an explicit transitivity rule.

Definition 3 (Type Equivalence) Two types T_1 and T_2 are said to be equivalent, denoted by $T_1 \equiv T_2$, if $T_1 \models T_2$ and $T_2 \models T_1$.

Under this definition, two types are equivalent if they define the same regular tree language. Interesting cases arise when we consider *distributivity* of types. Intuitively, $\{\text{int} \vee \text{null } f\} \equiv \{\text{int } f\} \vee \{\text{null } f\}$. However, care must be taken as, for example $[\text{int} \vee \text{null}] \not\equiv [\text{int}] \vee [\text{null}]$ since $[\text{int} \vee \text{null}] \models [1, \text{null}]$, but $[\text{int}] \vee [\text{null}] \not\models [1, \text{null}]$. As another example, consider:

```

define List1 as { int|null dat, null|List1 nxt }
define List2 as { null dat, null|List2 nxt }
define List3 as { int dat, null|List3 nxt }

```

Again, $\text{List1} \not\equiv \text{List2} \vee \text{List3}$ as, if $V = \{\text{dat} : \text{null}, \text{nxt} : \{\text{dat} : 1, \text{nxt} : \text{null}\}\}$, then $\text{List1} \models V$ but $\text{List2} \not\models V$ and $\text{List3} \not\models V$.

3.2 Subtyping

Whilst the above definitions characterise types in FW, they are not constructive from an implementation perspective. As our primary goal is to develop an efficient implementation, it is important to consider such algorithmic issues. For example, consider:

```

define Link as {int dt, LinkedList nxt}
define LinkedList as null | Link

```


One possible type for `LinkedList` is $\mu X.(\text{null} \vee \{\text{int } dt, X \text{ nxt}\})$, whilst another (equivalent) possibility is $\text{null} \vee \{\text{int } dt, \mu X.(\text{null} \vee \{\text{int } dt, X \text{ nxt}\}) \text{ nxt}\}$. Under Definition 3, recursive types are equivalent to their (infinite) set of unfoldings:

$$\begin{aligned} & \mu X.(\text{null} \vee \{\text{int } dt, X \text{ nxt}\}) \\ & \equiv (\text{null} \vee \{\text{int } dt, \mu X.(\text{null} \vee \{\text{int } dt, X \text{ nxt}\}) \text{ nxt}\}) \\ & \equiv (\text{null} \vee \{\text{int } dt, (\text{null} \vee \{\text{int } dt, \mu X.(\dots) \text{ nxt}\}) \text{ nxt}\}) \\ & \equiv \dots \end{aligned}$$

A critical question is: *how do we implement subtyping efficiently in the presence of such types?*

Amadio and Cardelli were the first to show that subtyping in the presence of recursive types was decidable [29]. Their system included function types, \top and \perp . Kozen *et al.* improved this by developing an $\mathcal{O}(n^2)$ algorithm [30]. The system presented here essentially extends this in a straightforward manner. Gapeyev *et al.* give an excellent overview of the subject [28] and, indeed, our subtype relation is identical to theirs, except for the inclusion of unions and other compound structures.

In a nominal type system, types correspond to *trees* and, thus, the subtype operator can be defined using rules such as:

$$\frac{T_1 \leq T_2}{[T_1] \leq [T_2]}$$

Here, a strong property holds that the “height” of T_1 is strictly less than $[T_1]$ — leading to a simple proof of termination since every type has finite height. In a structural type system, like FW, types correspond to graphs not trees. Defining the subtype operator using rules such as above leads to non-termination in the presence of cycles. To resolve this we employ ideas from co-induction [28].

The subtyping rules for FW are given in Figure 2. These employ judgements of the form “ $T_1 \leq T_2 \mid \mathcal{C}$ ”, which are read as: T_1 is a subtype of T_2 under assumptions \mathcal{C} . Essentially, the set of assumptions \mathcal{C} helps ensure the subtype operator terminates (we’ll return to this shortly). Ignoring the issue of assumption sets, the rules of Figure 2 are mostly straightforward. Since lists have value semantics, $[\text{int}] \leq [\text{any}]$ holds. Subtyping of records allows for *depth* and *width* [32]. Thus, $\{\text{int } x, \text{int } y\} \leq \{\text{any } x\}$ by S-REC.

Rule S-UNION3 captures distributivity over records. For example, under S-UNION3, it holds that $\{\text{int} \vee \text{null } x\} \leq \{\text{int } x\} \vee \{\text{null } x\}$. The following derivation illustrates a more complex example:

$$\begin{array}{ll} T_1 & = \{\text{int} \vee \text{null } x, \text{int} \vee \text{null } y\} \\ S_1 & = \{\text{int } x, \text{int } y\} \vee \{\text{null } x, \text{int } y\} \vee \{\text{int } x, \text{null } y\} \vee \{\text{null } x, \text{null } y\} \end{array}$$

1. $\{\text{int } x, \text{int } y\} \leq \{\text{int } x, \text{int } y\} \mid \{\dots\}$ (S-REFLEX)
2. $\{\text{null } x, \text{int } y\} \leq \{\text{null } x, \text{int } y\} \mid \{\dots\}$ (S-REFLEX)
3. $\{\text{int } x, \text{null } y\} \leq \{\text{int } x, \text{null } y\} \mid \{\dots\}$ (S-REFLEX)
4. $\{\text{null } x, \text{null } y\} \leq \{\text{null } x, \text{null } y\} \mid \{\dots\}$ (S-REFLEX)
5. $\{\text{int } x, \text{int } y\} \leq S_1 \mid \{\dots\}$ (S-UNION1, 1)
6. $\{\text{null } x, \text{int } y\} \leq S_1 \mid \{\dots\}$ (S-UNION1, 2)
7. $\{\text{int } x, \text{null } y\} \leq S_1 \mid \{\dots\}$ (S-UNION1, 3)
8. $\{\text{null } x, \text{null } y\} \leq S_1 \mid \{\dots\}$ (S-UNION1, 4)
9. $\{\text{int} \vee \text{null } x, \text{int} \vee \text{null } y\} \leq S_1 \mid \{\dots\}$ (S-UNION3, 5+6)
10. $\{\text{int} \vee \text{null } x, \text{null } y\} \leq S_1 \mid \{\dots\}$ (S-UNION3, 7+8)
11. $T_1 \leq S_1 \mid \emptyset$ (S-UNION3, 9+10)

Given the rules of Figure 2, we can give an algorithmic interpretation of subtyping in FW:

Definition 4 (Subtyping) *Let T_1 and T_2 be types. Then, T_1 is a subtype of T_2 , denoted $T_1 \leq T_2$, iff $T_1 \leq T_2 \mid \emptyset$.*

To show T_1 is a subtype of T_2 , we use the rules of Figure 2 starting with no assumptions. As we descend the type graph comparing components of T_1 and T_2 the set \mathcal{C} *always increases*. The S-INDUCT rule is critical here, as it terminates the recursion (by, essentially, treating the assumption set \mathcal{C} as a “visited” set). Furthermore, the size of \mathcal{C} can be bounded as follows: let m (resp. n) be the number of nodes in the type graph of T_1 (resp. T_2); then, every addition to \mathcal{C} made by a rule of Figure 2 corresponds to a pair (v, w) , where v and w are (respectively) nodes in the type graph of T_1 and T_2 — thus, $|\mathcal{C}|$ is $\mathcal{O}(m \cdot n)$. For example, consider the following type definition (where `LinkedList` is defined in §3.1):

```
define AnyList as null | {any dt, AnyList nxt}
```

```
AnyList f(LinkedList l):
  return l
```

For this function to be considered type safe, we need to show that $T_1 = \mu X.(\text{null} \vee \{\text{int dt}, X \text{ nxt}\})$ is a subtype of $S_1 = \mu X.(\text{null} \vee \{\text{any dt}, X \text{ nxt}\})$:

$$\begin{aligned} T_1 &= T_3 \vee T_2 \\ T_2 &= \{\text{int dt}, T_1 \text{ nxt}\} \\ T_3 &= \text{null} \\ S_1 &= S_3 \vee S_2 \\ S_2 &= \{\text{any dt}, S_1 \text{ nxt}\} \\ S_3 &= \text{null} \end{aligned}$$

$$\begin{aligned} 1. \quad & T_3 \leq S_3 \mid \{T_1 \leq S_1, \dots\} && \text{(S-REFLEX)} \\ 2. \quad & \text{int} \leq \text{any} \mid \{T_1 \leq S_1, \dots\} && \text{(S-ANY)} \\ 3. \quad & T_1 \leq S_1 \mid \{T_1 \leq S_1, \dots\} && \text{(S-INDUCT)} \\ 4. \quad & T_2 \leq S_2 \mid \{T_1 \leq S_1, \dots\} && \text{(S-REC, 2+3)} \\ 5. \quad & T_3 \leq S_1 \mid \{T_1 \leq S_1, \dots\} && \text{(S-UNION1, 1)} \\ 6. \quad & T_2 \leq S_1 \mid \{T_1 \leq S_1\} && \text{(S-UNION1, 4)} \\ 7. \quad & T_1 \leq S_1 \mid \emptyset && \text{(S-UNION2, 5+6)} \end{aligned}$$

Here, the assumption $T_1 \leq S_1$ is propagated through the proof, and terminates the recursion at (3) with the use of S-INDUCT.

Finally, the following derivation provides a useful sanity check that $\mu X. \{\text{int} \vee \text{null data}, \text{null} \vee X \text{ next}\}$ is *not* a subtype of $\mu X. \{\text{int data}, \text{null} \vee X \text{ next}\} \vee \mu X. \{\text{null data}, \text{null} \vee X \text{ next}\}$:

$$\begin{aligned} T_1 &= \{\text{int} \vee \text{null data}, \text{null} \vee T_1 \text{ next}\} \\ S_1 &= S_3 \vee S_2 \\ S_2 &= \{\text{int data}, \text{null} \vee S_2 \text{ next}\} \\ S_3 &= \{\text{null data}, \text{null} \vee S_3 \text{ next}\} \end{aligned}$$

$$\begin{aligned} 1. \quad & \text{int} \not\leq \text{null} \mid \{T_1 \leq S_1, \dots\} && \text{(By inspection of Fig 2)} \\ 2. \quad & \text{null} \not\leq \text{int} \mid \{T_1 \leq S_1, \dots\} && \text{(By inspection of Fig 2)} \\ 3. \quad & T_1 \not\leq \text{null} \mid \{T_1 \leq S_1, \dots\} && \text{(By inspection of Fig 2)} \\ 4. \quad & \text{int} \vee \text{null} \not\leq \text{null} \mid \{T_1 \leq S_1, \dots\} && \text{(S-UNION2, 1)} \\ 5. \quad & \text{int} \vee \text{null} \not\leq \text{int} \mid \{T_1 \leq S_1, \dots\} && \text{(S-UNION2, 2)} \\ 6. \quad & T_1 \not\leq S_3 \mid \{T_1 \leq S_1, \dots\} && \text{(S-REC, 4)} \\ 7. \quad & T_1 \not\leq S_2 \mid \{T_1 \leq S_1, \dots\} && \text{(S-REC, 5)} \\ 8. \quad & T_1 \not\leq \text{null} \vee S_3 \mid \{T_1 \leq S_1, \dots\} && \text{(S-UNION1, 3,6)} \\ 9. \quad & T_1 \not\leq \text{null} \vee S_2 \mid \{T_1 \leq S_1, \dots\} && \text{(S-UNION1, 3,7)} \\ 10. \quad & \text{null} \vee T_1 \not\leq \text{null} \vee S_3 \mid \{T_1 \leq S_1, \dots\} && \text{(S-UNION2, 8)} \\ 11. \quad & \text{null} \vee T_1 \not\leq \text{null} \vee S_2 \mid \{T_1 \leq S_1, \dots\} && \text{(S-UNION2, 9)} \\ 12. \quad & \{\text{int data}, \text{null} \vee T_1 \text{ next}\} \not\leq \{\text{null data}, \text{null} \vee S_3 \text{ next}\} \mid \{T_1 \leq S_1, \dots\} && \text{(S-REC, 10)} \\ 13. \quad & \{\text{int data}, \text{null} \vee T_1 \text{ next}\} \not\leq \{\text{int data}, \text{null} \vee S_2 \text{ next}\} \mid \{T_1 \leq S_1, \dots\} && \text{(S-REC, 11)} \\ 14. \quad & \{\text{int data}, \text{null} \vee T_1 \text{ next}\} \not\leq S_1 \mid \{T_1 \leq S_1\} && \text{(S-UNION1, 12,13)} \\ 15. \quad & T_1 \not\leq S_1 \mid \emptyset && \text{(S-UNION3, 14)} \end{aligned}$$

Note, in the above, S-UNION2 is used at (10) + (11); however, using S-UNION1 does not affect the outcome.

3.3 Soundness and Completeness

We now provide the necessary connections between the semantic and algorithmic interpretations. Unfortunately, performing a structural induction in the presence of assumption sets is challenging.

Definition 5 (Derivation) For a given subtyping computation $T_1 \leq T_2 \mid \mathcal{C}_1$, let \mathcal{D} denote its derivation. That is, the set of judgments of the form $T_3 \leq T_4 \mid \mathcal{C}_2$ generated by the rules of Figure 2.

Definition 6 (Projection) If \mathcal{D} is a derivation, then its projection is $\mathcal{D}^* = \{T_1 \leq T_2 \mid T_1 \leq T_2 \mid \mathcal{C} \in \mathcal{D}\}$.

Lemma 1 (Subtype Assumptions) Let T_1 and T_2 be types where computing $T_1 \leq T_2 \mid \emptyset$ generates an intermediate judgment $T_3 \leq T_4 \mid \mathcal{C}$ which is shown to hold. If $T_1 \leq T_2 \mid \emptyset$ then $T_3 \leq T_4 \mid \emptyset$.

Proof 1 Let \mathcal{D}_1 be the derivation for $T_1 \leq T_2 \downarrow \emptyset$ and \mathcal{D}_2 the derivation for $T_3 \leq T_4 \downarrow C$. By inspection of Figure 2, it follows that $\mathcal{D}_1^* \supseteq \mathcal{D}_2^*$. Let $C' = C - \{c\}$ for some arbitrary $c \in C$, and \mathcal{D}_3 be the derivation for $T_3 \leq T_4 \downarrow C'$. It follows that $\mathcal{D}_1^* \supseteq \mathcal{D}_3^* \supseteq \mathcal{D}_2^*$ as only S-INDUCT is affected by the difference C and C' . We can also conclude that $\mathcal{D}_3^* = T_3 \leq T_4 \downarrow C'$ was shown to hold. This is because assumptions do not interfere with the other rules from Figure 2. By an inductive argument, we can conclude $T_3 \leq T_4 \downarrow \emptyset$ is shown to hold. \square

Theorem 1 (Subtype Soundness) Let T and T' be types where $T \leq T'$. Then, $T \models T'$.

Proof 2 By induction on the structure of T and T' . The induction hypothesis is that if $T_1 \leq T_2 \downarrow \emptyset$ holds for some substructure of T and (resp.) T' , then $T_1 \models T_2$. Each case corresponds to a rule from Figure 2, although S-REFLEX, S-VOID, S-ANY are ignored since they follow immediately from Definition 1.

- Case $[T_1] \leq [T_2]$: By S-LIST, we have $T_1 \leq T_2 \downarrow C$ and, hence, $T_1 \leq T_2 \downarrow \emptyset$ by Lemma 1. Thus, it follows from Definition 1 that $[T_1] \models [T_2]$.
- Case $\{T_1 n_1, \dots, T_m n_m\} \leq \{T_1 T_1, \dots, T_n T_n\}$ where $n < m$: By S-REC, we have $T_i \leq S_i \downarrow C$ for $1 \leq i \leq m$ and, hence, $T_i \leq S_i \downarrow \emptyset$ by Lemma 1. This follows Definition 1 where only the first n fields of a record are tested for acceptance. Thus, $\{T_1 n_1, \dots, T_m n_m\} \models \{T_1 n_1, \dots, T_n n_n\}$.
- Case $T_1 \leq T_2 \vee T_3$: By S-UNION1, we have $T_1 \leq T_2 \downarrow C$ or $T_1 \leq T_3 \downarrow C$ and, hence, $T_1 \leq T_2 \downarrow \emptyset$ or $T_1 \leq T_3 \downarrow \emptyset$ by Lemma 1. Then $T_1 \models T_2$ or $T_1 \models T_3$. Therefore, $T_1 \models T_2 \vee T_3$ under Definition 1.
- Case $T_2 \vee T_3 \leq T_1$: By S-UNION2, we have $T_2 \leq T_1 \downarrow C$ and $T_3 \leq T_1 \downarrow C$ and, hence, $T_2 \leq T_1 \downarrow \emptyset$ and $T_3 \leq T_1 \downarrow \emptyset$ by Lemma 1. Then $T_2 \models T_1$ and $T_3 \models T_1$. Therefore, $T_2 \vee T_3 \models T_1$ under Definition 1.
- Case $\{T_1 f_1, \dots, T_1 \vee T'_1 f_1, \dots, T_n f_n\} \leq T'$: By S-UNION3 and Lemma 1, $\{T_1 f_1, \dots, T_i f_i, \dots, T_n f_n\} \leq T' \downarrow \emptyset$ and $\{T_1 f_1, \dots, T'_i f_i, \dots, T_n f_n\} \leq T' \downarrow \emptyset$. Then, $\{T_1 f_1, \dots, T_i \vee T'_i f_i, \dots, T_n f_n\} \models T'$ under Definition 1.
- Case $\overline{T} \rightarrow T \leq \overline{S} \rightarrow S$: By S-FUN, we have $T \leq S \downarrow C$ and $\overline{S} \leq \overline{T} \downarrow C$ and, hence, $T \leq S \downarrow \emptyset$ and $\overline{S} \leq \overline{T} \downarrow \emptyset$ by Lemma 1. It follows from Definition 1 that $T \models S$ and $\overline{S} \models \overline{T}$. Assume $\overline{T} \rightarrow T \models \overline{V} \rightarrow \overline{V}$. By Definition 1, we have that $\overline{V} \models \overline{T}$. Then, it follows that $\overline{V} \models S$. Hence, $\overline{S} \rightarrow S \models \overline{V} \rightarrow \overline{V}$. \square

Theorem 2 (Subtype Completeness) Let T and T' be types where $T \models T'$. Then, $T \leq T'$.

Proof 3 By induction on the structure of T and T' . The induction hypothesis is that if $T_1 \models T_2$ holds for some substructure of T and (resp.) T' , then $T_1 \leq T_2 \downarrow \emptyset$.


- Case $\text{void} \models T'$: Straightforward since $\text{void} \leq T'$ by S-VOID.
- Case $\text{any} \models T'$: By Definition 1, any accepts every value. Therefore, T' must accept every value and, hence, $T' = \text{any}$. Thus, $\text{any} \leq \text{any}$ by S-REFLEX.
- Case $\text{null} \models T'$: By Definition 1, null can only accept value null. Therefore, either $T' \equiv \text{null}$ or $T' \equiv \text{null} \vee T_1$ (for some T_1). Then, $T \leq T'$ by S-NULL and/or S-UNION1.
- Case $[T_1] \models [T_2]$: By induction hypothesis, $T_1 \leq T_2$ and, hence, $T \leq T'$ by S-LIST and S-UNION1.
- Case $\overline{T} \rightarrow T \models \overline{S} \rightarrow S$: By Definition 1, $T \models S$ and, by inductive hypothesis, $T \leq S$. It follows from Definition 1 that $\forall \overline{V}. [\overline{V} \models T_1 \implies \overline{V} \models S_1]$ (for $1 \leq i \leq n$) and, hence, $\overline{S} \models \overline{T}$. Finally, $\overline{T} \geq \overline{S}$ by inductive hypothesis and, hence, $T \leq T'$ by S-FUN.
- Case $\{\overline{T} n\} \models \{\overline{S} m\}$: By Definition 1, $m \leq n$. By induction hypothesis, $\overline{T} \leq \overline{S}$ and, hence, $T \leq T'$ by S-REC.
- Case $T_1 \vee T_2 \models T'$: By Definition 1, $T_1 \models T'$ and $T_2 \models T'$. By induction hypothesis, $T_1 \leq T'$ and $T_2 \leq T'$. Then, $T \leq T'$ by S-UNION2.
- Case $[T_1] \models T_2 \vee T_3$: If $T \models T_2$ or $T \models T_3$, then $T \leq T_2 \vee T_3$ by S-UNION1. Assume (for contradiction) that $T \not\models T_2$ and $T \not\models T_3$. It follows from Definition 1 that $T_2 = [T_4]$, $T_3 = [T_5]$ and, hence, $T_1 \models T_4 \vee T_5$. This gives the contradiction as $[T_1]$ will accept a list containing values from both T_4 and T_5 which neither $[T_4]$ nor $[T_5]$ could accept.
- Case $\overline{S} \rightarrow S \models T_1 \vee T_2$: If $T \models T_1$ or $T \models T_2$, then $T \leq T_1 \vee T_2$ by S-UNION1. Assume (for contradiction) that $T \not\models T_1$ and $T \not\models T_2$. Let $T_1 = \overline{W} \rightarrow W$ and $T_2 = \overline{U} \rightarrow U$. Then, the contradiction arises in the same manner as for lists. For example, consider return values. It follows from Definition 1 that $S \models W \vee U$. Then, T will accept a function with return values in both W and U which neither T_1 nor T_2 could accept.
- Case $\{\overline{S} f\} \models T_1 \vee T_2$: If $T \models T_1$ or $T \models T_2$, then $T \leq T_1 \vee T_2$ by S-UNION1. Assume $T \not\models T_1$ and $T \not\models T_2$. Then, some $S_i = T_3 \vee T_4$ must exist. Let $W = \{\overline{W} f\}$ where $W_i = T_3$ and $\forall_{j \neq i}. [W_j = S_j]$. Construct W' similarly, except with $W'_i = T_4$. It follows immediately that $W \models T_1 \vee T_2$ and $W' \models T_1 \vee T_2$. Hence, $T \leq T'$ by the inductive hypothesis and S-UNION3. \square

3.4 Representation

Another important question, from an implementation perspective, is how types can be represented efficiently. In particular, we desire a compact representation where equivalence testing is efficient¹. Our starting point is to provide a *graph-based* definition of types:

Definition 7 (Type) A type is a triple (V, E, K) , where V is a set of vertices, E is a set of directed, labelled edges of the form $v \xrightarrow{\ell} w$, and K gives the kind of each vertex (i.e. *NULL*, *INT*, *LIST*, *RECORD*, *UNION*, etc).


In this definition, the labels ℓ correspond either to field names or ϵ (i.e. no label). Thus, the recursive type $\mu X.(\text{null} \vee \{\text{int dt}, X \text{ nxt}\})$ and its first unfolding can be viewed as follows:



The graph on the right represents a single unfolding of that on the left. Here, circular nodes represent unions, square nodes represent records, and **NULL** is indicated by ground (for brevity, we omit field *dt*).

Definition 8 (Minimal Form) Let $T\downarrow$ denote a minimal form of a type T such that, for any $T' \equiv T$, it follows that $|T\downarrow| \leq |T'|$.

To compute $T\downarrow$ we can apply standard algorithms for minimising finite automata [49, 50, 51]. Observe that, for any given type, there may exist multiple minimal forms [52]. The following illustrates:



$$\begin{aligned} T_1 &= (\{1, 2, 3\}, \{1 \xrightarrow{\epsilon} 2, 1 \xrightarrow{\epsilon} 3, 3 \xrightarrow{\text{nxt}} 1\}, \{1 \mapsto \text{UNION}, 2 \mapsto \text{NULL}, 3 \mapsto \text{RECORD}\}) \\ T_2 &= (\{1, 2, 3\}, \{2 \xrightarrow{\epsilon} 3, 2 \xrightarrow{\epsilon} 1, 1 \xrightarrow{\text{nxt}} 2\}, \{1 \mapsto \text{RECORD}, 2 \mapsto \text{UNION}, 3 \mapsto \text{NULL}\}) \end{aligned}$$

Whilst it is clear that $T_1 \equiv T_2$, it is also the case that $T_1 \neq T_2$. Thus, minimising a type does not, by itself, guarantee a canonical form.

Definition 9 (Canonical Form) Let \hat{T} denote the canonical form of a type T such that, for any T_1 and T_2 where $T_1 \equiv T_2$, it follows that $\hat{T}_1 = \hat{T}_2$.

This notion of a canonical representation is similar to that found in the study of graph isomorphism, where *canonical labellings* are used by the most efficient algorithms for determining isomorphism [53, 54, 55, 56]. To compute the canonical form of a type T we first minimise it, and then determine a canonical labelling. Intuitively, the idea behind a canonical labelling is to ensure the labels given to vertices are always the same for equivalent types. To get a rough first approximation, we can sort vertices by type. For example, sorting according to $\text{NULL} < \text{RECORD} < \text{UNION}$ gives the following canonical representation of T_1 and T_2 :

$$T_3 = (\{1, 2, 3\}, \{3 \xrightarrow{\epsilon} 1, 3 \xrightarrow{\epsilon} 2, 2 \xrightarrow{\text{nxt}} 3\}, \{1 \mapsto \text{NULL}, 2 \mapsto \text{RECORD}, 3 \mapsto \text{UNION}\})$$

For simplicity, we will not discuss efficient algorithms for computing the canonical representation of a type. However, we will demonstrate that it exists and is always computable by a reduction to coloured graph isomorphism:

Definition 10 (Coloured Digraph) A coloured digraph is a triple (V, E, C) where V is a set of vertices, E a set of directed edges of the form $v \rightarrow w$ and C maps every vertex to a colour taken from $\{1, \dots, |V|\}$.

¹This stems from the fact that FW's typing algorithm employs fixed-point iteration and, in practice, this generates a lot of equivalence tests (see §4.4)

Lemma 2 (Canonical Form) For any type T , a canonical form \hat{T} can be computed such that, for any T_1 and T_2 where $T_1 \equiv T_2$, it follows that $\hat{T}_1 = \hat{T}_2$.

Proof 4 By reduction to coloured graph isomorphism. W.L.O.G. assume T is minimised. Construct a coloured graph (V_C, E_C, C_C) from a type (V_T, E_T, K_T) as follows:

- Let $V_C = V_T \cup \{\ell_{vw} \mid v \xrightarrow{\ell} w \in E_T \wedge \ell \neq \epsilon\}$.
- Let $E_C = \{v \rightarrow w \mid v \xrightarrow{\epsilon} w \in E_T\} \cup \{v \rightarrow \ell_{vw}, \ell_{vw} \rightarrow w \mid v \xrightarrow{\ell} w \in E_T\}$.
- Let $C_C = \{v \mapsto c(k) \mid K_T(v) = k \wedge v \in V_T\} \cup \{\ell_{vw} \mapsto c(\ell) \mid v \xrightarrow{\ell} w \in E_T\}$ where $c(\cdot)$ maps kinds and labels to unique colours.

Then, a canonical labelling of (V_C, E_C, C_C) can be computed with well-known algorithms that use backtracking search to find a minimum lexicographic ordering of edges (e.g. [53, 54]). \square

3.5 Other Operations

An important requirement is that our subtype relation forms a complete lattice. For any two types T_1 and T_2 , the least upper bound and greatest lower bound must be well-defined:

Definition 11 The least upper bound of two types T_1 and T_2 , denoted $T_1 \sqcup T_2$, is the least type T_3 where $T_1 \leq T_3$ and $T_2 \leq T_3$.

Definition 12 The greatest lower bound of two types T_1 and T_2 , denoted $T_1 \sqcap T_2$, is the greatest type T_3 where $T_3 \leq T_1$ and $T_3 \leq T_2$.

The least upper bound operation is needed to determine variable types at control-flow joins. Computing $T_1 \sqcup T_2$ is fairly straightforward as, by construction, it is equivalent to $T_1 \vee T_2$. The purpose of the greatest lower bound operation may be less apparent. Consider:

```
define MyType as int | [int] | [[int]]

int f(MyType x):
  if x is [any]:
    return |x|
  else:
    return x
```

The most precise type for x on the true branch is $[int] \vee [[int]]$. To update a variable's type after a type test, we *intersect* its type before the test with that used in the test. Or, in other words, we compute their greatest lower bound:

$$([int] \vee [[int]]) \sqcap [any] \implies [int] \vee [[int]]$$

Computing $T_1 \sqcap T_2$ is only slightly more involved than for $T_1 \sqcup T_2$. Recalling the correspondence between types and tree automata from §3.1, then $T_1 \sqcap T_2$ corresponds to the *intersection* of the tree automata described by T_1 and that described by T_2 . Again, standard algorithms for computing this are known [52].

Another question is how the type of x on the false branch is determined. We use a third operation on types called *greatest difference* (analogous to set difference):

Definition 13 The greatest difference of T_1 and T_2 , denoted $T_1 - T_2$, is the greatest type T_3 where $T_3 \leq T_1$ and $T_2 \sqcap T_3 = \text{void}$.

Using the greatest difference operation, we obtain the following type for x on the false branch:

$$([int] \vee [[int]]) - [any] \implies int$$

From this it follows that the last statement of function $f()$ above is type safe since, at this point, x has type **int**. Unfortunately, Definition 13 can only be approximated in FW. For example:

```
int f(any x):
  if x is int:
    ...
  else:
    ...
```

On the false branch, we compute the type $\text{any} - \text{int}$ for x ; however, this yields **any**, since there is no way to encode anything more precise in our language of types. In practice, this does not seem a major concern, as such types are not expressible at the source level.

4 Syntax, Semantics and Typing

Figure 3 gives the syntax for FW. As usual, the overbar (as in e.g. \bar{S}) is taken to indicate a list with appropriate separator(s). For simplicity, indentation is not explicitly described in the syntax, but is instead assumed. The notion of a runtime value (given by v) is very similar, but not identical, to that from §3.1. In fact, the only difference lies in the way function values are represented. Finally, for simplicity, we treat direct and indirect invocation with an identical statement. This means direct invocation corresponds to indirect invocation on a constant function value.

4.1 Semantics

A big-step operational semantics for FW is given for expressions in Figure 4, and a small-step semantics for statements in Figure 5. Here, Δ is the *runtime environment*, whilst v, w, u denote *runtime values*. A runtime environment Δ provides access to local variables in the current stack frame, whilst $\text{push}()$ and $\text{pop}()$ let us (resp.) create and destroy stack frames. Likewise, $\text{funDecl}(f)$ gives the declaration of a function f (we do not consider overloading). Finally, $\text{halt}(v)$ is taken to mean that the machine has halted, producing value v .

Since function invocation is via a statement, rather than an expression, it follows that expressions have *finite executions* (i.e. they cannot loop indefinitely). This means we can safely employ a big-step semantics for expressions, which simplifies the system (we return to discuss this later). In contrast, we must employ a small-step semantics for statements, as these can loop infinitely (i.e. through **while** loops or recursion).

4.2 Type Environments

Typing functions requires separate environments for each program point. We denote the environment for some point ℓ by Γ^ℓ , which captures the types of all variables immediately before the statement at that point. For example:

```
int f(int x) :
  y = x + 11
  return y2
```

The typing environments for this function are:

$$\begin{aligned}\Gamma^1 &= \{x \mapsto \text{int}\} \\ \Gamma^2 &= \{x \mapsto \text{int}, y \mapsto \text{int}\}\end{aligned}$$

Since y is undefined before line 1, it is absent from Γ^1 . To define how environments are combined at join points in the control-flow graph, we introduce a corresponding subtype relation:

Definition 14 Let Γ^{ℓ_1} and Γ^{ℓ_2} be typing environments. Then, $\Gamma^{\ell_1} \leq \Gamma^{\ell_2}$ iff $\forall v \in \text{dom}(\Gamma^{\ell_2}). [\Gamma^{\ell_1}[v] \leq \Gamma^{\ell_2}[v]]$.

For example, the following orderings are trivially valid:

$$\begin{aligned}\{v \mapsto \text{int}\} &\leq \{v \mapsto \text{any}\} \\ \{v \mapsto [\text{void}], w \mapsto \text{int}\} &\leq \{v \mapsto [\text{int}]\}\end{aligned}$$

On the other hand, the following do not hold:

$$\begin{aligned}\{v \mapsto [\text{int}]\} &\not\leq \{v \mapsto \text{int}\} \\ \{v \mapsto \text{int}\} &\not\leq \{v \mapsto \text{void}\}\end{aligned}$$

The following illustrates how type environments are combined:

```
int | [int] f(int x) :
  if y > 01 :
    y = [x]2
    z = 23
  else :
    y = []4
  return y5
```

The question is, what type does y have in Γ^5 ? We know that y has $[\text{int}]$ or $[\text{void}]$ type, and we desire the *most precise environment* capturing this. This corresponds to the *least upper bound* of the type environments involved:

$$\begin{aligned}\Gamma^5 &= \{x \mapsto \text{int}, y \mapsto [\text{int}], z \mapsto \text{int}\} \sqcup \{x \mapsto \text{int}, y \mapsto [\text{void}]\} \\ &\leftrightarrow \{x \mapsto \text{int}, y \mapsto [\text{int}]\}\end{aligned}$$

Here, $\Gamma^5[y] = [\text{int}]$ as $[\text{int}] \vee [\text{void}] \equiv [\text{int}]$. Furthermore, z cannot be included in Γ^5 as it is not defined on both branches.

Syntax:	
$F ::= T f(\overline{Tn}) : S$	
$S ::= n = e \mid n.f = e \mid n[e_1] = e_2 \mid n = e(\overline{e}) \mid \text{return } e \mid \text{if } n \text{ is } T : S \mid \text{else } S$	
$\quad \mid \text{while } e_1 < e_2 : S \mid S_1 ; S_2$	
$e ::= v \mid n \mid e.f \mid e_1[e_2] \mid [\overline{e}] \mid \{ \overline{n} : \overline{e} \}$	
$v ::= \text{null} \mid i \mid [\overline{v}] \mid \{ \overline{n} : \overline{v} \} \mid \&f$	

Figure 3: Syntax for Featherweight Whiley. Here, n represents variable identifiers, whilst $i \in \mathcal{I}$ represent integer constants.

Expression Reductions:		
$\frac{e \rightsquigarrow v}{\{ \overline{n} : e \} \rightsquigarrow \{ \overline{n} : v \}}$	$\frac{e \rightsquigarrow v}{[\overline{e}] \rightsquigarrow [\overline{v}]}$	(R-RECORD, R-LIST)
$\frac{e_2 \rightsquigarrow i \quad 0 \leq i \leq n}{e_1 \rightsquigarrow [v_0, \dots, v_n] \rightsquigarrow v_i}$	$\frac{e_2 \rightsquigarrow i \quad (i < 0 \vee i > n)}{e_1 \rightsquigarrow [v_0, \dots, v_n] \rightsquigarrow \text{err}}$	(R-LACCESS1, R-LACCESS2)
$\frac{e \rightsquigarrow \{ \overline{n} : \overline{v} \} \quad n_i = f}{e.f \rightsquigarrow v_i}$		(R-FACCESS)

Figure 4: Semantics for expressions in FW.

Statement Reductions:	
$\frac{\Delta \vdash e \rightsquigarrow v}{\langle \Delta, n = e ; S \rangle \longrightarrow \langle \Delta[n \mapsto v], S \rangle}$	(R-VASSIGN)
$\frac{\Delta \vdash e \rightsquigarrow v \quad \Delta(\overline{n}) = \{ \overline{n} : \overline{v} \} \quad v' = \{ \overline{n} : \overline{v} \} [f \mapsto v]}{\langle \Delta, n.f = e ; S \rangle \longrightarrow \langle \Delta[n \mapsto v'], S \rangle}$	(R-FASSIGN)
$\frac{\Delta(\overline{n}) = [u_0, \dots, u_n] \quad \Delta \vdash e_2 \rightsquigarrow v_1 \quad \Delta \vdash e_1 \rightsquigarrow i : \text{int} \quad 0 \leq i \leq n \quad v_2 = [u_0, \dots, u_{i-1}, v_1, u_{i+1}, \dots, u_n]}{\langle \Delta, n[e_1] = e_2 ; S \rangle \longrightarrow \langle \Delta[n \mapsto v_2], S \rangle}$	(R-LASSIGN)
$\frac{\Delta_1(v) \in T}{\langle \Delta_1, \text{if } v \text{ is } T : \{S_1\} \text{ else } \{S_2\} ; S_3 \rangle \longrightarrow \langle \Delta_2, S_1 ; S_3 \rangle}$	(R-IF1)
$\frac{\Delta_1(v) \notin T}{\langle \Delta_1, \text{if } v \text{ is } T : \{S_1\} \text{ else } \{S_2\} ; S_3 \rangle \longrightarrow \langle \Delta_2, S_2 ; S_3 \rangle}$	(R-IF2)
$\frac{\Delta \vdash e_1 \rightsquigarrow v_1 : \text{int} \quad \Delta \vdash e_2 \rightsquigarrow v_2 : \text{int} \quad v_1 \leq v_2}{\langle \Delta, \text{while } e_1 \leq e_2 : \{S_1\} ; S_2 \rangle \longrightarrow \langle \Delta, S_1 ; \text{while } e_1 \leq e_2 : \{S_1\} ; S_2 \rangle}$	(R-WHILE1)
$\frac{\Delta \vdash e_1 \rightsquigarrow v_1 : \text{int} \quad \Delta \vdash e_2 \rightsquigarrow v_2 : \text{int} \quad v_1 \not\leq v_2}{\langle \Delta, \text{while } e_1 \leq e_2 : \{S_1\} ; S_2 \rangle \longrightarrow \langle \Delta, S_2 \rangle}$	(R-WHILE2)
$\frac{\Delta \vdash e \rightsquigarrow \&f \quad e \rightsquigarrow v \quad \text{funDecl}(f) = T f(\overline{Tn}) : S_2 \quad \Delta_2 = \text{push}(\Delta_1, S_1, n)}{\langle \Delta_1, n = e(\overline{e}) ; S_1 \rangle \longrightarrow \langle \Delta_2[\overline{n} \mapsto \overline{v}], S_2 \rangle}$	(R-INVOKE)
$\frac{\Delta_1 \vdash e \rightsquigarrow v \quad \Delta_2, S_2, n = \text{pop}(\Delta_1)}{\langle \Delta_1, \text{return } e ; S_1 \rangle \longrightarrow \langle \Delta_2[n \mapsto v], S_2 \rangle}$	(R-RETURN)
$\frac{\Delta \vdash e \rightsquigarrow v \quad \epsilon = \text{pop}(\Delta)}{\langle \Delta, \text{return } e ; S \rangle \longrightarrow \text{halt}(v)}$	(R-HALT)

Figure 5: Small-step operational semantics for statements in FW.

Expression Typing:			
$\frac{}{\Gamma \vdash \text{null} : \text{null}}$	$\frac{i \in \mathbb{I}}{\Gamma \vdash i : \text{int}}$	(T-NULL, T-INT)	
$\frac{\text{funType}(f) = \bar{T} \rightarrow T}{\vdash \&f : \bar{T} \rightarrow T}$	$\frac{\{x \mapsto T\} \in \Gamma}{\Gamma \vdash x : T}$	(T-FUN, T-VAR)	
$\frac{\Gamma \vdash e_0 : T_0, \dots, \Gamma \vdash e_n : T_n}{\Gamma \vdash \{\bar{n} : \bar{e}\} : \{\bar{T} \bar{n}\}}$			(T-RECORD)
$\frac{\Gamma \vdash e_0 : T_0, \dots, \Gamma \vdash e_n : T_n}{\Gamma \vdash [\bar{e}] : [T_0 \vee \dots \vee T_n]}$			(T-LIST)
$\frac{\Gamma \vdash e_1 : [T] \quad \Gamma \vdash e_2 : \text{int}}{\Gamma \vdash e_1[e_2] : T}$			(T-LACCESS)
$\frac{\Gamma \vdash e : T \quad \text{ert}(T) = \{\dots, T_1 f, \dots\}}{\Gamma \vdash e.f : T_1}$			(T-FACCESS)
Effective Record Type:			
$\overline{\text{ert}(\{\bar{T} \bar{f}\})} = \{\bar{T} \bar{f}\}$			
$T'_1 = \text{ert}(T_1) = \{\bar{T} \bar{f}\} \quad T'_2 = \text{ert}(T_2) = \{\bar{T}' \bar{f}'\}$			
$F = \mathbf{dom}(T'_1) \cap \mathbf{dom}(T'_2)$			
$T_3 = \{T \bar{f} \mid \bar{f} \in F \wedge T = T'_1[\bar{f}] \vee T'_2[\bar{f}]\}$			
$\text{ert}(T_1 \vee T_2) = T_3$			

Figure 6: Typing rules for expressions in Featherweight Whieley.

4.3 Typing Expressions

The typing rules for expressions are given in Figure 6. These are presented using judgements of the form $\Gamma \vdash e : T$, which are taken to mean that under typing environment Γ , expression e has type T . Here, $\text{funType}(f)$ gives the static type of function f .

The *effective record type (ERT)* for type T is given by $\text{ert}(T)$. This is necessary for handling unions of records which have fields in common. The following illustrates:

$$\begin{aligned}
& \text{ert}\left(\{\text{int } x, [\text{int}] \text{ dt}\} \vee \{\text{any } x\}\right) = \{\text{any } x\} \\
& \text{ert}\left(\{\text{int } x\} \vee \{[\text{int}] x, \text{int } \text{cd}\}\right) = \{(\text{int} \vee [\text{int}]) x\} \\
& \text{ert}\left(\{\text{int } x, [\text{int}] \text{ dt}\} \vee [\text{int}]\right) = \{\} \\
& \text{ert}\left(\mu X. \{\text{int } x, (X \vee \text{null}) n\}\right) \\
& = \text{ert}\left(\{\text{int } x, (\mu X. \{\text{int } x, (X \vee \text{null}) n\} \vee \text{null}) n\}\right) \\
& = \{\text{int } x, (\mu X. \{\text{int } x, \dots\} \vee \text{null}) n\}
\end{aligned}$$

For a union of records, the ERT is formed from those fields common to all by taking the LUB of their types. This identifies those fields (if any) which may be accessed on such a type.

Finally, rule T-LIST from Figure 6 indicates the type of a list constructor is the LUB of all element types:

$$\begin{aligned}
& \vdash [1, [1]] \quad : \quad [\text{int} \vee [\text{int}]] \\
& \vdash [[], [1]] \quad : \quad [[\text{void}] \sqcup [\text{int}]] \implies [[\text{int}]]
\end{aligned}$$

Using the LUB of the element types ensures we always get the most precise type for a list constructor.

4.4 Typing Statements

As expected, typing statements in FW takes a somewhat different approach than normal. When typing a statement, we need to describe its *effect* on the typing environment. The typing rules are given in Figure 7. These rules are presented as judgements of the form $\Gamma \vdash S : \Gamma'$. Here, Γ represents the typing environment immediately before S , whilst Γ' represents that which holds immediately after. Thus, the effect of statement S is captured in the difference between Γ and Γ' . To clarify this, consider the simplest example:


```

int f(string x, string y):
  x = 11
  return x2

```

$\Gamma^1 = \{x \mapsto \text{string}, y \mapsto \text{string}\}$ gives the environment immediately before the assignment. Applying T-VASSIGN from Figure 7 yields the typing environment immediately after it, namely $\Gamma^2 = \{x \mapsto \text{int}, y \mapsto \text{string}\}$. Thus, the type of x is updated from `string` to `int` by the assignment.

Rule T-FASSIGN from Figure 7 makes use of the function `upert`(T_1, f, T_2), which returns an updated version of T_1 where field f now how has type T_2 . For unions of records this updates the field across all elements:

```

define LL as {int dt, LL nxt} | null

```

```

int f(LL link):
  if link != null1:
    link.dt = [1]2
    return 03
  ...

```

The typing environments determined for this function are:

$$\begin{aligned}
\Gamma^1 &= \{\text{link} \mapsto \mu X.(\{\text{int dt}, X \text{next}\} \vee \text{null})\} \\
\Gamma^2 &= \{\text{link} \mapsto \{\text{int dt}, \mu X.(\{\text{int dt}, X \text{next}\} \vee \text{null}) \text{next}\}\} \\
\Gamma^3 &= \{\text{link} \mapsto \{[\text{int}] \text{dt}, \mu X.(\{\text{int dt}, X \text{next}\} \vee \text{null}) \text{next}\}\}
\end{aligned}$$

Here, the `null` test removes the possibility of `link` being `null` in the true branch. Then, the assignment updates the type of field `dt` from `int` to `[int]`.

Rule T-LASSIGN also behaves somewhat unexpectedly:

```

[[int | [int]] f([int] xs):
  if |xs| > 1:
    xs[0] = [1]
  return xs

```

Variable `xs` has type `[int \vee [int]]` after the assignment as it now contains an element of `[int]` type. Of course, there are limitations on how well the type system can reason about list updates:

```

[[int]] f(int y):
  xs = [1]
  xs[0] = [1]
  return xs

```

Whilst `xs` must hold a value of type `[[int]]` at the return statement, the type system cannot prove this. The type of `xs` at this point is (conservatively) determined as `[int \vee [int]]` and, hence, this program is rejected.

Rule T-RETURN uses `thisFunType` (which holds the enclosing function’s type) to check the returned expression satisfies the required return type. The environment determined after a `return` statement is \perp (the subtype of all typing environments). Consider the following:

```

int f(int x):
  if x == 01:
    return 02
  else:
    y = x + 13
    return y4

```

The typing environment which holds before “`return y`” is:

$$\Gamma^4 = \perp \sqcup \{x \mapsto \text{int}, y \mapsto \text{int}\} \implies \{x \mapsto \text{int}, y \mapsto \text{int}\}$$

The special environment \perp is needed to ensure y retains its type at this point, since y is not defined on the true branch.

Finally, rule T-WHILE must determine the (least) fixed-point of the typing environment which holds within the body (as is common for flow-sensitive program analysis). This is more complicated than usual as, in some cases, fixed-point iteration is not guaranteed to terminate — we return to discuss this in the following section. This rule is also one reason we’re interested in efficient equivalence testing of types (recall §3.4). This is because testing for the fixed-point generates a lot of type equivalence tests.

Statement Typing:	
$\frac{\Gamma_0 \vdash S_1 : \Gamma_1 \quad \Gamma_1 \vdash S_2 : \Gamma_2}{\Gamma_0 \vdash S_1; S_2 : \Gamma_2}$	(T-SEQ)
$\frac{\Gamma \vdash e : T}{\Gamma \vdash n = e : \Gamma[n \mapsto T]}$	(T-VASSIGN)
$\frac{\Gamma \vdash n : T_1 \quad \Gamma \vdash e : T_2 \quad \text{ert}(T_1) = \{\overline{T f}\}}{\Gamma \vdash n.f = e : \Gamma[n \mapsto \text{upert}(T_1, f, T_2)]}$	(T-FASSIGN)
$\frac{\Gamma \vdash n : [T_1] \quad \Gamma \vdash e_1 : \text{int} \quad \Gamma \vdash e_2 : T_2}{\Gamma \vdash n[e_1] = e_2 : \Gamma[n \mapsto [T_1 \vee T_2]]}$	(T-LASSIGN)
$\frac{\Gamma \vdash e : \overline{T} \rightarrow T}{\Gamma \vdash e : T' \quad T' \leq T}$ $\frac{}{\Gamma \vdash n = e(\overline{\theta}) : \Gamma[n \mapsto T]}$	(T-INVOKE)
$\frac{\Gamma \vdash e : T \quad T \leq T_r \quad \text{thisFunType} = \overline{T} \rightarrow T_r}{\Gamma \vdash \text{return } e : \perp}$	(T-RETURN)
$\frac{\Gamma_0 \vdash v : T_1 \quad \Gamma_3 = \Gamma_1 \sqcup \Gamma_2 \quad \Gamma_0[v \mapsto T_1 \sqcap T_2] \vdash S : \Gamma_1 \quad \Gamma_0[v \mapsto T_1 - T_2] \vdash S' : \Gamma_2}{\Gamma_0 \vdash \text{if } v \text{ is } T_2 : \{S\} \text{ else } \{S'\} : \Gamma_3}$	(T-IF)
$\frac{\Gamma_0 \sqcup \Gamma_1 \vdash S : \Gamma_1 \quad \Gamma_0 \sqcup \Gamma_1 \vdash e_1 : \text{int} \quad \Gamma_0 \sqcup \Gamma_1 \vdash e_2 : \text{int}}{\Gamma_0 \vdash \text{while } e_1 \leq e_2 \{S\} : \Gamma_0 \sqcup \Gamma_1}$	(T-WHILE)
Update Effective Record Type:	
$\frac{T_1 = \{\overline{T f}\}}{\text{upert}(T_1, f_j, T_2) = T_1[f_j \mapsto T_2]}$	
$\frac{T_4 = \text{upert}(T_1, f, T_3) \vee \text{upert}(T_2, f, T_3)}{\text{upert}(T_1 \vee T_2, f, T_3) = T_4}$	

Figure 7: Typing rules for statements in FW.

5 Termination and Soundness

In this section, we prove two important properties for FW, namely: *soundness* and *termination*. Proving termination of FW's typing algorithm depends on showing the type environments $\Gamma_0, \dots, \Gamma_n$ for rule T-WHILE are monotonically increasing and, hence, always reach a fixed point. Unfortunately, by itself, this is not enough as the lattice has infinite height — however, we show how widening can be used to ensure termination without any loss of precision. Soundness requires that every well-typed program, when given appropriate input, will execute without getting stuck.

5.1 Termination

Showing termination of the typing algorithm relies on showing that we have a join-semi lattice (i.e. of the subtype relation) and that the transfer functions (i.e. the statement typing rules) are monotonic. However, these properties alone are insufficient to show termination of the typing rule T-WHILE from Figure 7. This is because the lattice of types has infinite height and, hence, the fixed-point algorithm may ascend it forever. To address this, we identify the special cases where it may happen, and demonstrate that widening can be employed to ensure termination *without loss of precision*. That is, assuming widening is performed at the correct point, the resulting typing for any program *will be as precise as is possible*.

As an example, consider the following simple FW function:

```

define LoopyList as [int | LoopyList]

LoopyList loopy(int n):
  x = [0]1
  while |x| < n:
    x[0] = x2
  return x

```

This method will cause a naive fix-point implementation of rule T-WHILE from Figure 7 to iterate forever. The following illustrates the environments that would be produced:

$$\begin{aligned}
\Gamma^1 &= \{n \mapsto \text{int}, x \mapsto [\text{int}]\} \\
\Gamma^2 &= \{n \mapsto \text{int}, x \mapsto [\text{int}]\} \\
\Gamma^2 &= \{n \mapsto \text{int}, x \mapsto [\text{int} \vee [\text{int}]]\} \\
\Gamma^2 &= \{n \mapsto \text{int}, x \mapsto [\text{int} \vee [\text{int} \vee [\text{int}]]]\} \\
\Gamma^2 &= \{n \mapsto \text{int}, x \mapsto [\text{int} \vee [\text{int} \vee [\text{int} \vee [\text{int}]]]]\} \\
&\dots
\end{aligned}$$

We can see that the type of x is essentially growing on every iteration of T-WHILE. Informally, we refer to this behaviour as being *divergent* (with respect to typing). Now, the crucial observation is that we can, in fact, give a precise type for x in Γ^2 , namely $\mu X.[\text{int} \vee X]$. Thus, the algorithm must infer this type from the divergent behaviour, as follows:

$$\begin{aligned}
\Gamma^2 &= \{n \mapsto \text{int}, x \mapsto [\text{int} \vee [\text{int} \vee [\text{int}]]]\} \\
\Gamma^2 &= \{n \mapsto \text{int}, x \mapsto [\text{int} \vee [\text{int} \vee [\text{int} \vee [\text{int}]]]]\} \\
&\dots \\
\Gamma^2 &= \{n \mapsto \text{int}, x \mapsto \mu X.[\text{int} \vee X]\} \\
\Gamma^2 &= \{n \mapsto \text{int}, x \mapsto [\text{int} \vee \mu X.[\text{int} \vee X]]\}
\end{aligned}$$

At this point, the computation terminates because $\mu X.[\text{int} \vee X] \equiv [\text{int} \vee \mu X.[\text{int} \vee X]]$ under Definition 3. The technique used here is often referred to as *widening* in the program analysis literature (see e.g. [57, 58, 59]) and, generally speaking, results in lost precision. However, in our setting, *widening does not lose precision* (although it is necessary to ensure termination).

Lemma 3 (Monotonicity) *Let $\Gamma_1 \vdash S : \Gamma_2$ and $\Gamma'_1 \vdash S : \Gamma'_2$ be typing environments. If $\Gamma_1 \leq \Gamma'_1$, then $\Gamma_2 \leq \Gamma'_2$.*

Proof 5 *Straightforward by inspection of Figure 7.* \square

Definition 15 (Divergence) *Let $\Gamma_1 \Longrightarrow S$ denote a diverging typing problem where $\Gamma_1 \vdash S : \Gamma_2, \Gamma_2 \vdash S : \Gamma_3, \dots$ is an infinite sequence such that $\forall i. [\Gamma_i < \Gamma_{i+1}]$.*

Lemma 4 (Divergence) *For a diverging typing problem $\Gamma_1 \Longrightarrow S$, a variable x exists where $\Gamma_i(x) < \Gamma_{i+1}(x)$ occurs infinitely often.*

Proof 6 *Straightforward as, otherwise, computation reaches a fixed point in finitely many steps as we have finitely many variables.* \square

Definition 16 *Let $T_1 \llbracket T_2 \rrbracket$ indicate type T_1 contains T_2 as one (or more) subcomponents.*

Definition 17 (Induction Variable) *A diverging typing problem $\Gamma_1 \Longrightarrow S$ is said to diverge on a variable x if $\Gamma_i(x) \equiv T \llbracket \Gamma_j(x) \rrbracket$ for $i > j$ occurs infinitely often.*

Lemma 5 (Induction Variable) *Every diverging typing problem $\Gamma_1 \Longrightarrow S$ is diverging upon at least one variable x .*

Proof 7 *Let x_i denote those variables where $\Gamma_i(x_i) < \Gamma_{i+1}(x_i)$ occurs infinitely often (Lemma 4). Then, there are infinitely many types generated for each x_i . Furthermore, new types can only be generated from constants, control-flow merges, type tests and through statements of the form $x_i[e_2] = e_1$ and $x_i.f = e_1$. Since the program structure is fixed, the divergence must stem from one or more statements of the form $x_i[e_1] = e_2$ and $x_i.f = e$. An infinite number of types can only be generated from such statements through divergence on x_i .* \square

An interesting question is whether or not we can strengthen our notion of an induction variable to say $\Gamma_i(x) \equiv T[\Gamma_{i+1}(x)]$ occurs infinitely often. We have found neither counter-example nor proof, but conjecture this is the case. However, one certainly cannot say that $\Gamma_i(x) \equiv T[\Gamma_{i+1}(x)]$ for all $i \geq 1$, as the following illustrates:

```
x = [0]
y = 1
while ... :
  x[0] = y
  y = x
```

Here, $\Gamma_1(x) = \Gamma_2(x) = [\text{int}]$. Furthermore, Lemma 5 might appear to fall down on this example:

```
x = [[0]]
while ... :
  x[0][0] = x[0]
```

This is because the key property that $\Gamma_{i+1}(x) = T[\Gamma_i(x)]$ does not appear to hold. However, the above example is not valid FW code! Rather, it's translation would be:

```
x = [[0]]
while ... :
  tmp = x[0]
  tmp[0] = x[0]
  x[0] = tmp
```

Thus, in this case, it follows that $\Gamma_2(\text{tmp}) = [\text{int} \vee \Gamma_1(\text{tmp})]$, $\Gamma_3(\text{tmp}) = [\text{int} \vee \Gamma_1(\text{tmp}) \vee \Gamma_2(\text{tmp})]$, etc.

Definition 18 (Least Widening) *Let $\Gamma_1 \Longrightarrow S$ be a typing problem which diverges on x . Then, a widening of x is the least type T such that either $\Gamma_1[x \mapsto T] \vdash S : \Gamma_2$ or $\Gamma_1[x \mapsto T] \Longrightarrow S$ no longer diverges on x .*

Lemma 6 (Least Widening) *For a typing problem $\Gamma_1 \Longrightarrow S$ which diverges on x , there exists a least widening of x .*

Proof 8 *Following Lemma 5 we have $\Gamma_i(x) \equiv T[\Gamma_j(x)]$ for $i > j$ infinitely often. Pick the smallest such i and j . Then, by construction, the least widening of x is $\mu X. \Gamma_i(x)[\Gamma_j(x) \mapsto X]$. \square*

Definition 18 indicates that a diverging typing problem may diverge on multiple independent variables. Widening just one of them does not necessarily eliminate all divergent behaviour. However, we can simply widen each individually until none remains.

We now return to consider the issue of showing that the least fixed point of the typing equations defined in Figure 7 is computable. The essential idea is that, for every non-divergent problem, we can just compute the fixed-point iteratively; but, for a divergent problem we immediately widen the diverging variable(s). The termination theorem states that a least solution to a (possibly diverging) typing problem $\Gamma_1 \vdash S : \Gamma_2$ can be computed:

Theorem 3 (Termination) *Given a typing environment Γ_1 , and a statement S , the least solution Γ_2 which satisfies $\Gamma_1 \vdash S : \Gamma_2$ is computable.*

Proof 9 *Straightforward given Lemma 3, Lemma 4 and Lemma 6. \square*

5.2 Soundness

In the following Lemmas, we use the term ‘‘closed expression’’ to represent an expression without free variables. Recall from §4.1 that our semantics for expressions given in Figure 4 guarantees finite execution traces. That is, expressions cannot execute indefinitely and will either produce a value, or halt with an error (indicated by **err**):

Lemma 7 (Finite Traces) *Let e_1 be an arbitrary closed expression. Then, either $e_1 \rightsquigarrow v$, or $e_1 \rightsquigarrow \text{err}$.*

Proof 10 *Straightforward by inspection of Figure 4. \square*

Essentially, Lemma 7 justifies our use of a big-step operational semantics for expressions, as it means we do not need to consider the case of non-terminating expressions (which is known to be difficult with a big-step operational semantics, and one important reason for using small-step [60]). As a result, our type-safety proof for expressions does not employ the usual notions of *progress* and *preservation*. Instead, we can prove type-safety directly as follows:

Lemma 8 (Effective Record Type) Let $\text{ert}(\mathbb{T}) = \{\overline{\mathbb{T} \mathbf{f}}\}$ for some type \mathbb{T} . Then $\mathbb{T} \leq \{\overline{\mathbb{T} \mathbf{f}}\}$.

Proof 11 Looking at Figure 6 there are two rules defining $\text{ert}(\mathbb{T})$. The case for $\text{ert}(\{\overline{\mathbb{T} \mathbf{n}}\})$ follows immediately, whilst $\text{ert}(\mathbb{T}_1 \vee \mathbb{T}_2)$ follows if we observe that $\mathbb{T}_1 \leq \mathbb{T}'_1$ and $\mathbb{T}_2 \leq \mathbb{T}'_2$ (by induction), and $\{\mathbb{T}'_1, \mathbb{T}'_2\} \leq \mathbb{T}_3$. This latter because \mathbb{T}_3 includes only fields present in both, and determines the type of each field \mathbf{f} from the least upper bound of its type in \mathbb{T}'_1 and \mathbb{T}'_2 . \square

Lemma 9 (Expression Safety) Let e be a well-typed, closed expression (i.e. $\vdash e : \mathbb{T}_e$). Then, $e \rightsquigarrow v$ where $\vdash v : \mathbb{T}_v$ and $\mathbb{T}_v \leq \mathbb{T}_e$.

Proof 12 By induction on the structure of $e : \mathbb{T}_e$. The *T-NUL*L, *T-INT* and *T-REAL* cases are immediate, since e in these cases is a value.

- Case “ $\{\overline{\mathbf{f} : e_{\mathbf{f}}}\}$ ”:
- Induction hypothesis guarantees, for each field \mathbf{f} with initialiser $e_{\mathbf{f}} : \mathbb{T}_{\mathbf{f}}$, that $e_{\mathbf{f}} \rightsquigarrow v_{\mathbf{f}}$ where $\vdash v_{\mathbf{f}} : \mathbb{T}'_{\mathbf{f}}$ and $\mathbb{T}'_{\mathbf{f}} \leq \mathbb{T}_{\mathbf{f}}$. Thus, $\vdash \{\overline{\mathbf{f} : e_{\mathbf{f}}}\} \rightsquigarrow \{\overline{\mathbf{f} : v_{\mathbf{f}}}\} : \mathbb{T}_v$ and so $\mathbb{T}_v \leq \mathbb{T}_e$ since we have $\{\mathbb{T}'_{\mathbf{f}} \mathbf{f}\} \leq \{\mathbb{T}_{\mathbf{f}} \mathbf{f}\}$ by *S-REC*.
- Case “ $[\overline{e}]$ ”:
- Induction hypothesis guarantees, for each element expression $e_i : \mathbb{T}_i$, that $e_i \rightsquigarrow v_i$ where $\vdash v_i : \mathbb{T}'_i$ and $\mathbb{T}'_i \leq \mathbb{T}_i$. Thus, $\mathbb{T}_v \leq \mathbb{T}_e$ because $\mathbb{T}'_1 \sqcup, \dots, \sqcup \mathbb{T}'_n \leq \mathbb{T}_1 \sqcup, \dots, \sqcup \mathbb{T}_n$ by Definition 11.
- Case “ $e_1[e_2]$ ”:
- Induction hypothesis guarantees, for source expression $e_1 : \mathbb{T}_e$, that $e_1 \rightsquigarrow [\overline{v}]$ where $\vdash [\overline{v}] : [\mathbb{T}']$ and $[\mathbb{T}_v] \leq [\mathbb{T}_e]$. Likewise, for the index expression, it guarantees $e_2 \rightsquigarrow i$ where $\vdash i : \text{int}$. Thus, assuming the list access is not out-of-bounds, we have $\mathbb{T}_v \leq \mathbb{T}_e$.
- Case “ $e_1.\mathbf{f}$ ”:
- Induction hypothesis guarantees, for source expression $e_1 : \mathbb{T}_1$, that $e_1 \rightsquigarrow v_1$ where $\vdash v_1 : \{\mathbb{T}'_{\mathbf{f}} \mathbf{f}\}$ and $\{\mathbb{T}'_{\mathbf{f}} \mathbf{f}\} \leq \mathbb{T}_1$. *T-FACCESS* also ensures that $\text{ert}(\mathbb{T}_1) = \{\overline{\mathbb{T}_{\mathbf{f}} \mathbf{f}}\}$ and, hence, it follows by Lemma 8 that $\{\mathbb{T}'_{\mathbf{f}} \mathbf{f}\} \leq \mathbb{T}_1 \leq \{\overline{\mathbb{T}_{\mathbf{f}} \mathbf{f}}\}$. Finally, by *S-REC*, we have $\{\mathbb{T}'_{\mathbf{f}} \mathbf{f}\}.\mathbf{f} \leq \{\overline{\mathbb{T}_{\mathbf{f}} \mathbf{f}}\}.\mathbf{f}$ and, hence, $\mathbb{T}_v \leq \mathbb{T}_e$. \square

At this point, we can now extend our notion of expression evaluation to include variables from the runtime environment, as follows:

Definition 19 (Expression Evaluation) Let Δ be a runtime environment, and e an arbitrary expression with free variables given by $\text{freeVars}(e) = \overline{x}$. If $e[x \mapsto \Delta(x)] \rightsquigarrow v$ where $v : \mathbb{T}$, then we say e evaluates to a value v of type \mathbb{T} under Δ , denoted by $\Delta \vdash e \rightsquigarrow v : \mathbb{T}$.

The following notion of a *safe abstraction* captures the relationship between type environments and their corresponding runtime environments:

Definition 20 (Safe Abstraction) Let Γ be a typing environment and Δ a runtime environment. Then, Γ is a safe abstraction of Δ , denoted $\Gamma \approx \Delta$, iff $\text{dom}(\Gamma) \subseteq \text{dom}(\Delta)$ and, for all $x \in \text{dom}(\Gamma)$ where $\Gamma(x) = \mathbb{T}_x$, $\Delta(x) = v$ and $\vdash v : \mathbb{T}_v$, it holds that $\mathbb{T}_v \leq \mathbb{T}_x$.

Observe that we cannot require $\text{dom}(\Gamma) = \text{dom}(\Delta)$, as might be expected, since runtime environments are the product of actual execution paths. Consider an *if* statement with a variable x defined only on one branch. After the statement, $x \notin \Gamma$ since x was not defined on both branches. However, if execution had proceeded through the branch where x was defined, then we would have $x \in \Delta$.

Lemma 10 (Safe Join) Let $\Gamma_1 \approx \Delta_1$. Then, $\Gamma_1 \sqcup \Gamma_2 \approx \Delta_1$ for any Γ_2 .

Proof 13 Straightforward, since by construction $\Gamma_1 \leq \Gamma_1 \sqcup \Gamma_2$ and, hence, $\forall x. [\Gamma_1(x) \leq (\Gamma_1 \sqcup \Gamma_2)(x)]$ holds by Definition 14. \square

Theorem 4 (Progress) Let Δ_1 be a runtime environment and Γ_1 a typing environment where $\Gamma_1 \approx \Delta_1$. If $\Gamma_1 \vdash S : \Gamma_2$, then either $\langle \Delta_1, S ; S' \rangle \longrightarrow \langle \Delta_2, S'' \rangle$ or $\langle \Delta_1, S ; S' \rangle \longrightarrow \text{halt}(v)$.

Proof 14 By induction on the structure of S .

- Case “ $x = e ; S'$ ”:
- Since $\Gamma_1 \vdash S : \Gamma_2$ and $\Gamma_1 \approx \Delta_1$, it follows by *T-VASSIGN* that $\Gamma_1 \vdash e : \mathbb{T}_e$ and, by Lemma 9, that $\Delta_1 \vdash e \rightsquigarrow v$. Hence, $\langle \Delta_1, x = e ; S' \rangle \longrightarrow \langle \Delta_1[x \mapsto v], S' \rangle$.

- Case “ $x.f = e ; S'$ ”: Since $\Gamma_1 \vdash S : \Gamma_2$ and $\Gamma_1 \approx \Delta_1$, it follows by *T-FASSIGN* that $\Gamma_1 \vdash x : T_x$ where $\text{ert}(T_x) = \{\overline{T}f\}$ and, hence, $\Delta_1 \vdash x \rightsquigarrow \{\overline{f} : \overline{v}\}$ by Lemma 9. Likewise, $\Gamma_1 \vdash e : T_e$ and, by Lemma 9, $\Delta_1 \vdash e \rightsquigarrow v_1$. Hence, $\langle \Delta_1, x.f = e ; S' \rangle \longrightarrow \langle \Delta_1[x \mapsto v_2], S' \rangle$, where $v_2 = \{\overline{f} : \overline{v}\}[f \mapsto v_1]$.
- Case “ $x[e_1] = e_2 ; S'$ ”: Since $\Gamma_1 \vdash S : \Gamma_2$ it follows by *T-LASSIGN* that $\Gamma_1 \vdash e_1 : \text{int}$, $\Gamma_1 \vdash e_2 : T_2$ and $\Gamma_1 \vdash x : [T_3]$. Thus, $\Gamma_1 \approx \Delta_1$ and Lemma 9 imply that $\Delta_1 \vdash e_1 \rightsquigarrow v_1 : \text{int}$, $\Delta_1 \vdash e_2 \rightsquigarrow v_2 : T_2$ and $\Delta_1 \vdash x \rightsquigarrow v_x$ (where $v_x = [\overline{v}]$). Thus, $\langle \Delta_1, x[e_1] = e_2 ; S' \rangle \longrightarrow \langle \Delta_1[x \mapsto v'_x], S' \rangle$, where $v'_x = v_x[v_1 \mapsto v_2]$.
- Case “ $\text{return } e ; S'$ ”: Since $\Gamma_1 \vdash S : \Gamma_2$ and $\Gamma_1 \approx \Delta_1$, it follows by *T-RETURN* that $\Gamma_1 \vdash e : T_e$ and, by Lemma 9, that $\Delta_1 \vdash e \rightsquigarrow v_1$. Hence, $\langle \Delta_1, \text{return } e ; S' \rangle \longrightarrow \langle \Delta_2, S'' \rangle$ or $\langle \Delta_1, \text{return } e \rangle \longrightarrow \text{halt}(v)$.
- Case “ $\text{if } v \rightsquigarrow T : S_1 \text{ else } : S_2$ ”: Since $\Gamma_1 \vdash S : \Gamma_2$ and $\Gamma_1 \approx \Delta_1$, it follows that $\Delta_1 \vdash x \rightsquigarrow v_x$. Thus, either *T-IF1* or *T-IF2* apply, leading to a transition of the form $\langle \Delta_1, S \rangle \longrightarrow \langle \Delta_1, S' \rangle$.
- Case “ $\text{while } e_1 \leq e_2 : S$ ”: Since $\Gamma_1 \vdash S : \Gamma_2$ it follows by *T-WHILE* that $\Gamma_1 \vdash e_1 : \text{int}$ and $\Gamma_2 \vdash e_2 : \text{int}$. Then, $\Gamma_1 \approx \Delta_1$ implies that $\Delta_1 \vdash e_1 \rightsquigarrow v_1 : \text{int}$ and $\Delta_1 \vdash e_2 \rightsquigarrow v_2 : \text{int}$ by Lemma 9. Thus, either *R-WHILE1* or *R-WHILE2* apply, leading to a transition of the form $\langle \Delta_1, S \rangle \longrightarrow \langle \Delta_2, S' \rangle$.
- Case “ $n = e(\overline{e}) ; S'$ ”: Since $\Gamma_1 \vdash S : \Gamma_2$ and $\Gamma_1 \approx \Delta_1$, it follows by *T-VOKE* that $\Delta_1 \vdash e : \overline{T} \rightarrow T$, and that $\overline{\Delta}_1 \vdash e : \overline{S}$ where $\overline{S} \leq \overline{T}$. Hence, by *R-VOKE* we have $\langle \Delta_1, n = e(\overline{e}) ; S' \rangle \longrightarrow \langle \Delta_2, S'' \rangle$.

□

Theorem 5 (Preservation) Let $\langle \Delta_1, S \rangle$ be a runtime environment and Γ_1 a typing environment where $\Gamma_1 \approx \Delta_1$. If $\Gamma_1 \vdash S : \Gamma_2$ for $S \notin \{n = e(\overline{e}), \text{return } e\}$ and $\langle \Delta_1, S \rangle \longrightarrow \langle \Delta_2, S' \rangle$, then $\Gamma_2 \approx \Delta_2$.

Proof 15 By induction on the structure of S .

- Case “ $x = e$ ”: By *T-VASSIGN* we have $\Gamma_2 = \Gamma_1[x \mapsto T_e]$, where $\Gamma_1 \vdash e : T_e$. Similarly, by *R-VASSIGN*, we have $\Delta_2 = \Delta_1[x \mapsto v]$ where $\Delta_1 \vdash e \rightsquigarrow v : T_v$. Since $\Gamma_1 \approx \Delta_1$, it follows that $T_v \leq T_e$ by Lemma 9 and, hence, that $\Gamma_2 \approx \Delta_2$.
- Case “ $x.f = e$ ”: By *T-FASSIGN* we have $\Gamma_1 \vdash e : T_e$, $\Gamma_1 \vdash x : T_x$ and $\Gamma_2 = \Gamma_1[x \mapsto T'_x]$, where T'_x is formed from T_x by updating field f to T_e . Similarly, by *R-VASSIGN*, we have $\Delta_2 = \Delta_1[x \mapsto v']$ where $\Delta_1 \vdash e \rightsquigarrow v : T_v$, $\Delta_1(x) = \{\overline{n} : \overline{v}\}$ and $v' = \{\overline{n} : \overline{v}\}[f \mapsto v]$. Since $T_v \leq T_e$ by Lemma 9, we have $\Gamma_2 \approx \Delta_2$.
- Case “ $x[e_1] = e_2$ ”: By *T-LASSIGN* we have $\Gamma_2 = \Gamma_1[x \mapsto [T_1 \sqcup T_2]]$, where $\Gamma_1 \vdash x : [T_1]$ and $\Gamma_1 \vdash e_2 : T_2$. By *R-LASSIGN*, $\Delta_1 \vdash e_1 \rightsquigarrow i : \text{int}$, $\Delta_1 \vdash e_2 \rightsquigarrow v : T'_2$, $\Delta_1 \vdash x \rightsquigarrow [\overline{v}] : [T'_1]$ and, since $\Gamma_1 \approx \Delta_1$, we have $T'_1 \leq T_1$ and $T'_2 \leq T_2$ by Lemma 9. Furthermore, $\Delta_2 = \Delta_1[x \mapsto [\overline{w}]]$ where $w_i = v$ and $\forall_{j \neq i}. [w_j = v_j]$. So, $\Delta_2 \vdash x \rightsquigarrow [\overline{w}] : [T'_1 \sqcup T'_2]$ and, thus, $\Gamma_2 \approx \Delta_2$ since $T'_1 \sqcup T'_2 \leq T_1 \sqcup T_2$.
- Case “ $\text{if } v \rightsquigarrow T : S_1 \text{ else } : S_2$ ”: By *T-IF* we have $\Gamma_1 \vdash S_1 : \Gamma_3$ and $\Gamma_1 \vdash S_2 : \Gamma_4$, where $\Gamma_2 = \Gamma_3 \sqcup \Gamma_4$. Likewise, we have either $\langle \Delta_1, S_1 \rangle \longrightarrow \langle \Delta_2, \epsilon \rangle$ by *R-IF1* and $\Gamma_3 \approx \Delta_2$ (by the induction hypothesis), or $\langle \Delta_1, S_2 \rangle \longrightarrow \langle \Delta_2, \epsilon \rangle$ by *R-IF2* and $\Gamma_4 \approx \Delta_2$. Thus, $\Gamma_3 \sqcup \Gamma_4 \approx \Delta_2$ by Lemma 10 and, hence, $\Gamma_2 \approx \Delta_2$ (since $\Gamma_2 = \Gamma_3 \sqcup \Gamma_4$).
- Case “ $\text{while } e_1 \leq e_2 : S$ ”: In executing this loop, the machine goes through a sequence of zero or more transitions of the form $\langle \Delta_1, S \rangle \longrightarrow \langle \Delta'_1, S \rangle, \langle \Delta'_1, S \rangle \longrightarrow \langle \Delta''_1, S \rangle$, etc. Considering *T-WHILE*, we arrive at a typing $\Gamma_1 \sqcup \Gamma_n \vdash S : \Gamma_1 \sqcup \Gamma_n$. Since $\Gamma_1 \approx \Delta_1$, we have by Lemma 10 that $\Gamma_1 \sqcup \Gamma_n \approx \Delta_1$. Then, by the induction hypothesis, we get that $\Gamma_1 \sqcup \Gamma_n \approx \Delta'_1, \Gamma_1 \sqcup \Gamma_n \approx \Delta''_1$ and so on.

□

Essentially, we limit our preservation theorem to the *intraprocedural* cases as typing is inherently an intraprocedural activity. A coherence theorem handles the movement between functions:

Definition 21 (Invocation) Let $\langle \Delta_1, S \rangle \Longrightarrow \langle \Delta_2, S' \rangle$ represent the shortest possible sequence which matches $\langle \Delta_1, S \rangle \longrightarrow \dots \longrightarrow \langle \Delta_2, S' \rangle$.

Theorem 6 (Coherence) Let $\langle \Delta_1, n = e(\overline{e}) ; S \rangle$ be a runtime environment and Γ_1 a typing environment where $\Gamma_1 \approx \Delta_1$. If $\Gamma_1 \vdash n = e(\overline{e}) : \Gamma_2$ and $\langle \Delta_1, n = e(\overline{e}) ; S \rangle \Longrightarrow \langle \Delta_2, S \rangle$, then $\Gamma_2 \approx \Delta_2$.

Proof 16 By induction on the structure of the induced call graph. Let k indicate the number of function invocations occurring in the sequence $\langle \Delta_1, n = e(\overline{e}) ; S \rangle \Longrightarrow \langle \Delta_2, S \rangle$:

- $k = 0$. Straightforward as, if no other function invocations occur, then $\Gamma_2 \approx \Delta_2$ essentially follows immediately from Theorem 5.
- $k = n$. Consider the first such subsequence $\langle \Delta_3, n = e(\bar{e}) ; S' \rangle \Longrightarrow \langle \Delta_4, S' \rangle$. If Γ_3 and Γ_4 are the typing environments before and after this statement, then $\Gamma_3 \approx \Delta_3$ by Theorem 5, and $\Gamma_4 \approx \Delta_4$ by the induction hypothesis. At this point, we have reduce the problem to the $k - 1$ case.

□

6 Related Work

In this section, we concentrate primarily on work relating to Whiley’s flow-sensitive type system.

The Strongtalk system pioneered the use of structural typing for describing structures from an untyped world (i.e. SmallTalk) [61]. It is perhaps the most closely aligned work to Whiley. The essential idea is succinctly captured in the following quote (from [61]):

“Smalltalk is an unusually flexible and expressive language. Any type system for Smalltalk should place a high priority on preserving its essential flavor.”

The paper hints that some form of recursive structural typing was supported, although few details are given and no formalisation is included. Unfortunately, structural subtyping was subsequently dropped from Strongtalk [62], in part due to the complexity of error messages produced. Although not the subject of this paper, we believe the quality of error messages in a structural type system can be improved by storing the names associated with types at their declarations. Then, when reporting an error, the compiler looks up the type in question (or those isomorphic to it) and matches any name(s) in the current file, or imported files.

The work of Guha *et al.* focuses on flow-sensitive type checking for JavaScript [63]. This assumes programmer annotations are given for parameters, and operates in two phases: first, a flow analysis inserts special runtime checks; second, a standard (i.e. flow-insensitive) type checker operates on the modified AST. The system retypes variables as a result of runtime type tests, although only simple forms are permitted. Recursive data types are not supported, although structural subtyping would be a natural fit here; furthermore, the system assumes sequential execution (true of JavaScript), since object fields can be retyped.

Tobin-Hochstadt and Felleisen consider the problem of typing previously untyped Racket (aka Scheme) programs and develop a technique called *occurrence typing* [64]. Their system will retype a variable within an expression dominated by a type test. Like Whiley, they employ union types to increase the range of possible values from the untyped world which can be described; however, they fall short of using full structural types for capturing arbitrary structure. Furthermore, in Racket, certain forms of aliasing are possible, and this restricts the points at which occurrence typing is applicable.

The earlier work of Aiken *et al.* is similar to that of Tobin-Hochstadt and Felleisen [13]. This operates on a function language with single-assignment semantics. They support more expressive types, but do not consider recursive structural types. Furthermore, instead of type checking directly on the AST, conditional set constraints are generated and solved. Following the soft typing discipline, their approach is to insert runtime checks at points which cannot be shown type safe.

The Java Bytecode Verifier requires a flow-sensitive type checker [65]. Since locals and stack locations are untyped in Java Bytecode, it must infer their types to ensure type safety. Like Whiley, the verifier updates the type of a variable after an assignment, and combines types at control-flow join points using a least upper bound operator. However, it does not update the type of a variable after an `instanceof` test. Furthermore, the Java class hierarchy does not form a join semi-lattice. To deal with this, the bytecode verifier uses a simplified least upper bound operator which ignores interfaces altogether, instead relying on runtime checks to catch type errors (see e.g. [66]).

Type qualifiers constrain the possible values a variable may hold. CQual is a flow-sensitive qualifier inference supporting numerous type qualifiers, including those for synchronisation and file I/O [17]. CQual does not account for the effects of conditionals and, hence, retyping is impossible. Fähndrich and Leino discuss a system for checking non-null qualifiers in the context of C# [34]. Here, variables are annotated with `NotNull` to indicate they cannot hold `null`. Non-null qualifiers are interesting because they require variables be retyped after conditionals (i.e. retyping `v` from `Nullable` to `NotNull` after `v != null`). Fähndrich and Leino hint at the use of retyping, but focus primarily on issues related to object constructors. Ekman *et al.* implemented this system within the JustAdd compiler, although few details are given regarding variable retyping [35]. Pominville *et al.* also briefly discuss a flow-sensitive non-null analysis built using SOOT, which does retype variables after `!=null` checks [33]. The JACK tool is similar, but focuses on bytecode verification instead [38]. JavaCOP provides an expressive language for writing type system extensions,

including non-null types [43]. This system is flow-insensitive and cannot account for the effects of conditionals; as a work around, the tool allows assignment from a nullable variable x to a non-null variable if this is the first statement after a $x \neq \text{null}$ conditional.

Information Flow Analysis is the problem of tracking the flow of information, usually to restrict certain flows for security reasons. The work of Hunt and Sands is relevant here, since they adopt a flow-sensitive approach [18]. Their system is presented in the context of a simple While language not dissimilar to ours, although they do not account for the effect of conditionals. Russo *et al.* use an extended version of this system to compare dynamic and static approaches [67]. They demonstrate that a purely dynamic system will reject programs that are considered type-safe under the Hunt and Sands system. JFlow extends Java with statically checked flow annotations which are flow-insensitive [68]. Finally, Chugh *et al.* developed a constraint-based (flow-insensitive) information flow analysis of JavaScript [69].

Typestate Analysis focuses on flow-sensitive reasoning about the state of objects, normally to enforce temporal safety properties. Typestates are finite-state automata which can encode usage rules for common APIs (e.g. a file is never read before being opened), and were pioneered by Strom and Yellin [70, 71]. Fink *et al.* present an interprocedural, flow-sensitive typestate verification system which is staged to reduce overhead [72]. Bodden *et al.* develop an interprocedural typestate analysis which is flow-sensitive at the intra-procedural level [73]. This is a hybrid system which attempts to eliminate all failure points statically, but uses dynamic checks when necessary. This was later extended to include a backward propagation step that improves precision [74].

6.1 Type Inference

Type inference is the process of inferring a type for each variable in an untyped (or partially untyped) program. This is generally done by generating and solving either unification- (e.g. [75]) or set-constraints (e.g. [76, 46, 77]) over the program in question.

Numerous type inference systems have been developed for object-oriented languages (e.g. [76, 78, 46, 77, 15, 16]). These, almost exclusively, assume the original program is completely untyped and employ set constraints (see [79]) as the mechanism for inferring types. As such, they address a different problem to that studied here. To perform type inference, such systems generate constraints from the program text, formulate them as a directed graph and solve them using an algorithm similar to transitive closure. When the entire program is untyped, type inference must proceed across method calls (known as *interprocedural analysis*).

Gagnon *et al.* present a technique for converting Java Bytecode into an intermediate representation with a single static type for each variable [80]. Key to this is the ability to infer static types for the local variables and stack locations used in the bytecode. Since local variables are untyped in Java bytecode, this is not always possible as they can — and often do — have different types at different points; in such situations, a variable is split as necessary into multiple variables each with a different type.

Bierman *et al.* formalise the type inference mechanism to be included in C# 3.0, the latest version of the C# language [9]. This uses a very different technique from us, known as *bidirectional type checking*, which was first developed for System F by Pierce and Turner [81]. This approach is suitable for C# 3.0 which does not permit variables to have different types at different program points.

7 Conclusion

We have presented the flow-sensitive and structural type system developed for the Whiley language. This permits variables to be declared implicitly, have multiple types within a function, and be retyped after runtime type tests. The semantics of Whiley also allows flexible typing of updates to compound structures. Additionally, subtyping between user-defined types is implicit. The result is a statically-typed language which, for the most part, has the look and feel of a dynamic language. An open source implementation of Whiley is available from <http://whiley.org>.

We formalised the type system using a core calculus called Featherweight Whiley (FW). We provided both a semantic and algorithmic interpretation of types including a coinductively defined subtype relation. The typing algorithm requires a distinct typing environment at each program point (much like flow-sensitive program analysis). Typing environments are merged at control-flow joins, whilst a greatest lower bound (resp. greatest difference) operator is used for retyping variables on the true (resp. false) branches of runtime type tests. We have sketched soundness and termination properties for this system — Full details can be found in [82].

Finally, there are several improvements to Whiley's type system that could be considered in the future. Most notably, the absence of generic types (i.e. type polymorphism) can result in the need for additional

runtime type tests, compared with dynamically typed languages. The greatest difference operator, as discussed in §3.5, often produces conservative results. This could be resolved by including negation and intersection types [46, 83].

Acknowledgements. The authors would especially like to thank Neal Glew, Nick Cameron, Roman Klapaukh, Art Protin and Lindsay Groves for helpful comments on earlier drafts of this paper.

References

- [1] R. Cartwright and M. Fagan. Soft typing. In *Proc. PLDI*, pages 278–292. ACM Press, 1991.
- [2] D. Ancona, M. Ancona, A. Cuni, and N. D. Matsakis. RPython: a step towards reconciling dynamically and statically typed OO languages. In *Proc. DLS*, pages 53–64. ACM Press, 2007.
- [3] John K. Ousterhout. Scripting: Higher-level programming for the 21st century. *IEEE Computer*, 31(3):23–30, 1998.
- [4] Diomidis Spinellis. Java makes scripting languages irrelevant? *IEEE Software*, 22(3):70–71, 2005.
- [5] Ronald Prescott Loui. In praise of scripting: Real programming pragmatism. *IEEE Computer*, 41(7):22–26, 2008.
- [6] B. Bloom, J. Field, N. Nystrom, J. Östlund, G. Richards, R. Strnisa, J. Vitek, and T. Wrigstad. Thorn: robust, concurrent, extensible scripting on the JVM. In *Proc. OOPSLA*, pages 117–136, 2009.
- [7] Linda Dailey Paulson. Developers shift to dynamic programming languages. *IEEE Computer*, 40(2):12–15, 2007.
- [8] The scala programming language. <http://lamp.epfl.ch/scala/>.
- [9] G. Bierman, E. Meijer, and M. Torgersen. Lost in translation: formalizing proposed extensions to C#. In *Proc. OOPSLA*, pages 479–498, 2007.
- [10] D. Remy and J. Vouillon. Objective ML: An effective object-oriented extension to ML. *TOPS*, 4(1):27–50, 1998.
- [11] J. Siek and W. Taha. Gradual typing for objects. In *Proc. ECOOP*, volume 4609 of *LNCS*, pages 151–175. Springer-Verlag, 2007.
- [12] T. Wrigstad, F. Nardelli, S. Lebresne, J. Östlund, and J. Vitek. Integrating typed and untyped code in a scripting language. In *Proc. POPL*, pages 377–388, 2010.
- [13] Alexander S. Aiken, Edward L. Wimmers, and T. K. Lakshman. Soft typing with conditional types. In *Proc. POPL*, pages 163–173, 1994.
- [14] Cormac Flanagan. Hybrid type checking. In *Proc. POPL*, pages 245–256. ACM Press, 2006.
- [15] C. Anderson, P. Giannini, and S. Drossopoulou. Towards type inference for javascript. In *Proc. ECOOP*, pages 428–452, 2005.
- [16] M. Furr, J.-H. An, J. Foster, and M. Hicks. Static type inference for Ruby. In *Proc. SAC*, pages 1859–1866. ACM Press, 2009.
- [17] Jeffrey S. Foster, Tachio Terauchi, and Alex Aiken. Flow-sensitive type qualifiers. In *Proc. PLDI*, pages 1–12. ACM Press, 2002.
- [18] Sebastian Hunt and David Sands. On flow-sensitive security types. In *Proc. POPL*, pages 79–90. ACM Press, 2006.
- [19] Jong-Deok Choi, Michael Burke, and Paul Carini. Efficient flow-sensitive interprocedural computation of pointer-induced aliases and side effects. In *Proc. POPL*, pages 232–245. ACM Press, 1993.
- [20] D. Pearce and J. Noble. Implementing a language with flow-sensitive and structural typing on the JVM. In *Proc. BYTECODE*, 2011.
- [21] D. Malayeri and J. Aldrich. Integrating nominal and structural subtyping. In *Proc. ECOOP*, pages 260–284, 2008.

- [22] Luca Cardelli. Structural subtyping and the notion of power type. In *Proc. POPL*, pages 70–79. ACM Press, 1988.
- [23] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley.
- [24] Nurudeen Lameed and Laurie J. Hendren. Staged static techniques to efficiently implement array copy semantics in a MATLAB JIT compiler. In *Proc. CC*, volume 6601 of *LNCS*, pages 22–41. Springer-Verlag, 2011.
- [25] Natarajan Shankar. Static analysis for safe destructive updates in a functional language. In *In Proc. LOPSTR*, pages 1–24, 2001.
- [26] Martin Odersky. How to make destructive updates less destructive. In *Proc. POPL*, pages 25–36, 1991.
- [27] Paul F. Dietz and Daniel D. Sleator. Two algorithms for maintaining order in a list. In *Proc. STOC*, pages 365–372. ACM Press, 1987.
- [28] Vladimir Gapeyev, Michael Y. Levin, and Benjamin C. Pierce. Recursive subtyping revealed. *JFP*, 12(6):511–548, 2002.
- [29] Roberto M. Amadio and Luca Cardelli. Subtyping recursive types. *ACM TOPLAS*, 15:575–631, 1993.
- [30] Dexter Kozen, Jens Palsberg, and Michael I. Schwartzbach. Efficient recursive subtyping. In *Proc. POPL*, pages 419–428, 1993.
- [31] Tony Hoare. Null references: The billion dollar mistake, presentation at qcon, 2009.
- [32] B. C. Pierce. *Types and Programming Languages*. MIT Press, 2002.
- [33] P. Pominville, F. Qian, R. Vallée-Rai, L. Hendren, and C. Verbrugge. A framework for optimizing Java using attributes. In *Proc. CC*, pages 334–554, 2001.
- [34] Manuel Fähndrich and K. Rustan M. Leino. Declaring and checking non-null types in an object-oriented language. In *Proc. OOPSLA*, pages 302–312. ACM Press, 2003.
- [35] Torbjörn Ekman and Görel Hedin. Pluggable checking and inferencing of non-null types for Java. *JOT*, 6(9):455–475, 2007.
- [36] M. Cielecki, J. Fulara, K. Jakubczyk, and L. Jancewicz. Propagation of JML non-null annotations in Java programs. In *Proc. PPPJ*, pages 135–140, 2006.
- [37] Patrice Chalin and Perry R. James. Non-null references by default in Java: Alleviating the nullity annotation burden. In *Proc. ECOOP*, pages 227–247. Springer, 2007.
- [38] C. Male, D. J. Pearce, A. Potanin, and C. Dymnikov. Java bytecode verification for @NonNull types. In *Proc. CC*, pages 229–244, 2008.
- [39] Laurent Hubert. A non-null annotation inferencer for java bytecode. In *Proc. PASTE*, pages 36–42. ACM, 2008.
- [40] L. Hubert, T. Jensen, and D. Pichardie. Semantic foundations and inference of non-null annotations. In *Proc. FMOODS*, pages 132–149, 2008.
- [41] F Barbanera and M Dezan-CianCaglini. Intersection and union types. In *Proc. of TACS*, volume 526 of *LNCS*, pages 651–674, 1991.
- [42] Atsushi Igarashi and Hideshi Nagira. Union types for object-oriented programming. *JOT*, 6(2), 2007.
- [43] C. Andreae, J. Noble, S. Markstrum, and T. Millstein. A framework for implementing pluggable type systems. In *Proc. OOPSLA*, pages 57–74. ACM Press, 2006.
- [44] ISO/IEC. international standard ISO/IEC 9899, programming languages — C, 1990.
- [45] A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. *TOPLAS*, 23(3):396–459, May 2001.
- [46] Alexander Aiken and Edward L. Wimmers. Type inclusion constraints and type inference. In *Proc. FPCA*, pages 31–41. ACM Press, 1993.

- [47] Flemming M. Damm. Subtyping with union types, intersection types and recursive types. volume 789 of *LNCS*, pages 687–706. 1994.
- [48] Alain Frisch, Giuseppe Castagna, and Véronique Benzaken. Semantic subtyping: Dealing set-theoretically with function, union, intersection, and negation types. *JACM*, 55(4):19:1–19:64, 2008.
- [49] J. E. Hopcroft. An $n \log n$ algorithm for minimizing the states in a finite-automaton. In *TMC*, pages 189–196. 1971.
- [50] Bruce W. Watson and Jan Daciuk. An efficient incremental dfa minimization algorithm. *NLE*, 9:49–64, 2003.
- [51] Marco Almeida, Nelma Moreira, and Rogério Reis. Incremental DFA minimisation. In *CIAA*, pages 39–48, 2010.
- [52] J. Hopcroft, R. Motwani, and J. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley, 2006.
- [53] B. D. McKay. Practical graph isomorphism. In *NMC*, volume 30, pages 45–87, 1981.
- [54] Babai and Luks. Canonical labeling of graphs. In *Proc. STOC*, pages 171–183, 1983.
- [55] T. Miyazaki. The complexity of mckay’s canonical labelling algorithm. *Groups and Computation, II*, 28:239–256, 1997.
- [56] J. Faulon. Isomorphism, automorphism partitioning, and canonical labeling can be solved in polynomial-time for molecular graphs. *JCICS*, 38(3):432–444, 1998.
- [57] François Bourdoncle. Efficient chaotic iteration strategies with widenings. In *Proceedings of the conference on Formal Methods in Programming and their Applications*, pages 128–141. Springer-Verlag, 1993.
- [58] Patrick Cousot and Radhia Cousot. Comparison of the Galois connection and widening/narrowing approaches to abstract interpretation. *BIGRE*, 74:107–110, 1991.
- [59] Flemming Nielson, Hanne R. Nielson, and Chris L. Hankin. *Principles of Program Analysis*. Springer-Verlag, 1999.
- [60] Xavier Leroy and Hervé Grall. Coinductive big-step operational semantics. *Information and Computation*, 207(2):284–304, 2009.
- [61] G. Bracha and D. Griswold. Strongtalk: Typechecking smalltalk in a production environment. In *OOPSLA*, pages 215–230, 1993.
- [62] Gilad Bracha. The strongtalk type system for smalltalk.
- [63] A. Guha, C. Saftoiu, and S. Krishnamurthi. Typing local control and state using flow analysis. In *Proc. ESOP*, pages 256–275, 2011.
- [64] Sam Tobin-Hochstadt and Matthias Felleisen. Logical types for untyped languages. In *Proc. ICFP*, pages 117–128, 2010.
- [65] Tim Lindholm and Frank Yellin. *The Java Virtual Machine Specification*. Addison Wesley, second edition, 1999.
- [66] X. Leroy. Java bytecode verification: algorithms and formalizations. *Journal of Automated Reasoning*, 30(3/4):235–269, 2003.
- [67] A. Russo and A. Sabelfeld. Dynamic vs. static flow-sensitive security analysis. In *Proc. CSF*, pages 186–199, 2010.
- [68] Andrew C. Myers. JFlow: Practical mostly-static information flow control. In *Proc. POPL*, pages 228–241, 1999.
- [69] R. Chugh, J. A. Meister, R. Jhala, and S. Lerner. Staged information flow for javascript. In *Proc. PLDI*, pages 50–62, 2009.
- [70] R. Strom and S. Yemini. Tpestate: A programming language concept for enhancing software reliability. *IEEE TSE*, 12(1):157–171, 1986.

- [71] Robert E. Strom and Daniel M. Yellin. Extending tpestate checking using conditional liveness analysis. *IEEE TSE*, 19(5):478–485, 1993.
- [72] S. Fink, E. Yahav, N. Dor, G. Ramalingam, and E. Geay. Effective tpestate verification in the presence of aliasing. *ACM TOSEM*, 17(2):1–9, 2008.
- [73] Eric Bodden, Patrick Lam, and Laurie J. Hendren. Finding programming errors earlier by evaluating runtime monitors ahead-of-time. In *Proc. ESEC/FSE*, pages 36–47. ACM Press, 2008.
- [74] E. Bodden. Efficient hybrid tpestate analysis by determining continuation-equivalent states. In *Proc. ICSE*, pages 5–14, 2010.
- [75] Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17:348–375, 1978.
- [76] Jens Palsberg and Michael I. Schwartzbach. Object-oriented type inference. In *Proc. OOPSLA*, pages 146–161. ACM Press, 1991.
- [77] T. Wang and S.F. Smith. Precise constraint-based type inference for Java. In *Proc. ECOOP*, pages 99–117. Springer-Verlag, 2001.
- [78] N. Oxhøj, J. Palsberg, and M. Schwartzbach. Making type inference practical. In *Proc. ECOOP*, pages 329–349. Springer-Verlag, 1992.
- [79] Alexander Aiken and Edward L. Wimmers. Solving systems of set constraints. In *Proceedings of LICS*, pages 329–340, 1992.
- [80] E. Gagnon, L. Hendren, and G. Marceau. Efficient inference of static types for java bytecode. In *Proc. SAS*, pages 199–219, 2000.
- [81] B. Pierce and D. Turner. Local type inference. *ACM TOPLAS*, 22(1):1–44, 2000.
- [82] D. J. Pearce and J. Noble. Flow-sensitive types for whiley. Technical Report ECSTR10-23, Victoria University of Wellington, 2010.
- [83] A. Frisch, G. Castagna, and V. Benzaken. Semantic subtyping. In *Proc. LICS*, pages 137–146. IEEE Computer Society Press, 2002.