

Sound and Complete Flow Typing with Unions, Intersections and Negations

David J. Pearce

Victoria University of Wellington
Wellington, New Zealand
{djp}@ecs.vuw.ac.nz

Abstract. Flow typing is becoming a popular mechanism for typing existing programs written in untyped languages (e.g. JavaScript, Racket, Groovy). Such systems require intersections for the true-branch of a type test, negations for the false-branch, and unions to capture the flow of information at meet points. Type systems involving unions, intersections and negations require a subtype operator which is non-trivial to implement. Frisch *et al.* demonstrated that this problem was decidable. However, their proof was not constructive and does not lend itself naturally to an implementation. In this paper, we present a sound and complete algorithm for subtype testing in the presence of unions, intersections and negations.

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 [3–6].

A common complaint against statically typed languages is the need for often unnecessarily verbose type declarations. Hindley-Milner Type inference [7, 8] is a common approach to addressing this problem, where type declarations are inferred automatically. Scala [9], C#3.0 [10] and OCaml [11] provide good examples of this in an imperative setting. However, such languages still require each program variable to have exactly one type. Flow typing offers an alternative to Hindley-Milner type inference where a variable may have different types at different program points. The technique is adopted from flow-sensitive program analysis and has been used for non-null types [12–15], information flow [16–18], purity checking [19] and more [12, 13, 20–26].

Few languages exist which incorporate flow typing directly. Typed Racket [23, 24] provides a *typed* sister language for *untyped* Racket, where flow typing is used to capture common idioms in the untyped language. Similarly, the recent 2.0 release of the popular Groovy language includes a flow typing algorithm [27]. Again, this is designed to handle common idioms in (previously) untyped Groovy programs. Finally, the Whyley language employs flow-typing to give it the look-and-feel of an untyped language [28–30].

1.1 Flow Typing

A defining characteristic of flow typing is the ability to *retype* a variable — that is, assign it a completely unrelated type. The JVM Bytecode Verifier [31], perhaps the most widely-used example of a flow typing system, provides a good illustration:

```
public static float convert(int):
    iload 0 // load register 0 on stack
    i2f     // convert int to float
    fstore 0 // store float to register 0
    fload 0 // load register 0 on stack
    freturn // return value on stack
```

In the above, register 0 contains the parameter value on entry and, initially, has type `int`. The type of register 0 is subsequently changed to `float` by the `fstore` bytecode. To ensure type safety, the JVM bytecode verifier employs a typing algorithm based upon dataflow analysis [32]. This tracks the type of a variable at each program point, allowing it easily to handle the above example.

Flow typing can also retype variables after conditionals. A *non-null type system* (e.g. [12–15]) prevents variables which may hold `null` from being dereferenced. The following illustrates:

```
int cmp(String s1, @NonNull String s2) {
    if (s1 != null) {
        return s1.compareTo(s2);
    } else {
        return -1;
    }
}
```

The modifier `@NonNull` indicates a variable definitely cannot hold `null` and, hence, that it can be safely dereferenced. To deal with the above example, a non-null type system will retype variable `s1` to `@NonNull` on the true branch — thus allowing it to type check the subsequent dereference of `s1`.

Whiley [28–30] employs a flow type system to give it the look-and-feel of a dynamically typed language. Variable retyping through conditionals is supported using the `is` operator (similar to `instanceof` in Java) as follows:

```
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`” 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 variable `s` to the appropriate `Shape` before accessing its fields.

1.2 Unions, Intersections and Negations

Union types (e.g. $T_1 \vee T_2$) are commonly used in flow typing systems to capture the type of variables at meet points. For example, consider this code snippet:

```
if ...:
    x = 1
else:
    x = true
...
```

After the assignment $x=1$, the type of variable x is `int`. Likewise, after the assignment $x=true$ it is `bool`. Finally, x has type $\text{int} \vee \text{bool}$ immediately after the `if` statement (i.e. at the meet point). This indicates x can hold either an `int` or a `bool` at that point.

Retyping variables after runtime type tests is typically achieved through a type system which supports both *intersections* (e.g. $T_1 \wedge T_2$) and *negations* (e.g. $\neg T_1$). For example, consider again this code snippet:

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

To determine the type of variable s on the true branch, we *intersect* its declared type (i.e. `Shape`) with the type test (i.e. `Circle`). Likewise, on the false branch, we compute the difference of these two types (i.e. `Shape - Circle`). Observe that the difference of two types in such a system is given by: $T_1 - T_2 \equiv T_1 \wedge \neg T_2$.

1.3 Contributions

Subtype testing (i.e. establishing whether $T_1 \leq T_2$ holds or not) is a challenging algorithmic problem for a type system involving unions, intersections and negations. In particular, we desire that subtyping is both *sound* and *complete* with respect to a semantic model where types are viewed as sets. The former requires $T_1 \leq T_2$ holds *only* when T_1 is a subset of T_2 , whilst the latter requires that $T_1 \leq T_2$ holds *whenever* T_1 is a subset of T_2 . Frisch *et al.* demonstrated that this problem was decidable [33]. However, their proof was not constructive and does not lend itself naturally to an implementation. In this paper, we present a sound and complete algorithm for subtyping in the presence of unions, intersections and negations. This contrasts with previous flow type systems (e.g. [23, 24]) which are shown sound, but not complete.

2 A Flow-Typing Calculus — FT

We now introduce our flow-typing calculus, FT, within which we frame our flow typing problem. The calculus is specifically kept to a minimum to allow us to succinctly capture the important issues. In this section, we introduce the syntax, semantics and subtyping rules for FT. We tacitly assume at this point that an appropriate subtyping operator exists. Subsequently, in §3 and §4, we will detail the algorithms which implement this operator (and which are the core contribution of this paper).

2.1 Types

The following gives a *syntactic* definition of types in FT:

$$T ::= \text{any} \mid \text{int} \mid (T_1, \dots, T_n) \mid \neg T \mid T_1 \wedge \dots \wedge T_n \mid T_1 \vee \dots \vee T_n$$

Here, `any` represents \top , `int` the set of all integers and (T_1, \dots, T_n) represents tuples with one or more elements. The union $T_1 \vee T_2$ is a type whose values are in T_1 *or* T_2 . Union types are generally useful in flow typing systems, as they can characterise types generated at meet points in the control-flow graph. The intersection $T_1 \wedge T_2$ is a type whose values are in T_1 *and* T_2 . Intersections are needed in our flow type system to capture the type of a variable (e.g. x) after a type test (e.g. x is T). The type $\neg T$ is the *negation* type containing those values *not* in T . Thus, $\neg \text{any}$ represents \perp (i.e. the empty set) and we will often write `void` as a short-hand for this. Negations are also useful for capturing the type of a variable on the false branch of a type test. Finally, we make some simplifying assumptions regarding unions and intersections: *namely, that elements are unordered and duplicates are removed*. Thus, $T_1 \vee T_2$ is indistinguishable from $T_2 \vee T_1$. Likewise, $T_1 \vee T_1$ is not distinguishable from T_1 . Whilst these simplifications are not strictly necessary, they simplify our presentation. Furthermore, they can be implemented easily enough by sorting elements according to a fixed total ordering of types.

To better understand the meaning of types in FT, it is helpful to give a *semantic interpretation* (following e.g. [34–36, 33]). The aim is to give a set-theoretic model where subtype corresponds to subset. The *domain* \mathbb{D} of values in our model consists of the integers and all records constructible from values in \mathbb{D} :

$$\mathbb{D} = \mathbb{Z} \cup \left\{ (v_1, \dots, v_n) \mid v_1 \in \mathbb{D}, \dots, v_n \in \mathbb{D} \right\}$$

Definition 1 (Type Semantics). *Every type T is characterised by the set of values it accepts, given by $\llbracket T \rrbracket$ and defined as follows:*

$$\begin{aligned} \llbracket \text{any} \rrbracket &= \mathbb{D} \\ \llbracket \text{int} \rrbracket &= \mathbb{Z} \\ \llbracket (T_1, \dots, T_n) \rrbracket &= \{ (v_1, \dots, v_n) \mid v_1 \in \llbracket T_1 \rrbracket, \dots, v_n \in \llbracket T_n \rrbracket \} \\ \llbracket \neg T \rrbracket &= \mathbb{D} - \llbracket T \rrbracket \\ \llbracket T_1 \wedge \dots \wedge T_n \rrbracket &= \llbracket T_1 \rrbracket \cap \dots \cap \llbracket T_n \rrbracket \\ \llbracket T_1 \vee \dots \vee T_n \rrbracket &= \llbracket T_1 \rrbracket \cup \dots \cup \llbracket T_n \rrbracket \end{aligned}$$

It is important to distinguish the *syntactic* representation from the *semantic* model of types. The former corresponds (roughly speaking) to a physical machine representation, whilst the latter is a mathematical ideal. As such, the syntactic representation diverges from the semantic model and, to compensate, we must establish a correlation between them. For example `int` and $\neg \neg \text{int}$ have distinct syntactic representations, but are semantically indistinguishable. Similarly for $(\text{int} \vee (\text{int}, \text{int}), \text{any})$ and $(\text{int}, \text{any}) \vee ((\text{int}, \text{int}), \text{any})$.

Ultimately, we want to construct a subtyping algorithm that is both *sound* and *complete* (i.e. that $T_1 \leq T_2 \iff \llbracket T_1 \rrbracket \subseteq \llbracket T_2 \rrbracket$). The distinction between syntactic and semantic forms presents a significant challenge in doing this.

Syntax:		Operational Semantics:	
$t ::=$	<i>terms:</i>	$\frac{\Delta \vdash t_k \longrightarrow t'_k}{\Delta \vdash (\dots, t_k, \dots) \longrightarrow (\dots, t'_k, \dots)}$	(E-TUP)
x	<i>variable</i>		
(t_1, \dots, t_n)	<i>tuple</i>		
$f t_1$	<i>application</i>	$\frac{\Delta \vdash t_1 \longrightarrow t'_1}{\Delta \vdash f t_1 \longrightarrow f t'_1}$	(E-APP1)
$f(T x) = t_1 \text{ in } t_2$	<i>declaration</i>		
$\text{if}(x \text{ is } T) t_1 \text{ else } t_2$	<i>type test</i>	$\frac{\Delta(f) = \langle T, x, t_2 \rangle \quad v_1 \in \llbracket T \rrbracket}{\Delta \vdash f v_1 \longrightarrow t_2[x \mapsto v_1]}$	(E-APP2)
$v ::=$	<i>values:</i>	$\frac{\Delta[f \mapsto \langle T, x, t_1 \rangle] \vdash t_2 \longrightarrow t'_2}{\Delta \vdash f(T x) = t_1 \text{ in } t_2 \longrightarrow f(T x) = t_1 \text{ in } t'_2}$	(E-DEC1)
i	<i>integer</i>	$\Delta \vdash f(T x) = t_1 \text{ in } v_2 \longrightarrow v_2$	(E-DEC2)
(v_1, \dots, v_n)	<i>tuple</i>	$\frac{v_1 \in \llbracket T \rrbracket}{\Delta \vdash \text{if}(v_1 \text{ is } T) t_2 \text{ else } t_3 \longrightarrow t_2}$	(E-IF1)
		$\frac{v_1 \notin \llbracket T \rrbracket}{\Delta \vdash \text{if}(v_1 \text{ is } T) t_2 \text{ else } t_3 \longrightarrow t_3}$	(E-IF2)

Fig. 1. Syntax and (small-step) operational semantics for FT.

2.2 Syntax & Semantics

Figure 1 gives the syntax of FT along with a small-step operational semantics, where $\Delta[f \mapsto \langle T, x, t \rangle]$ returns Δ with f now mapped to a triple $\langle T, x, t \rangle$ representing its declaration. Here, T denotes the parameter type, x the parameter name and t the function body. Similarly, $t[x \mapsto v]$ returns the term t with all occurrences of x now substituted with v . To avoid issues of variable capture, we assume parameter names are unique and may only occur within their function body (i.e. that for $f(T x) = t_1 \text{ in } t_2$ parameter x can only occur in t_1).

From the figure, we see that a semantic notion of type is explicitly required for the operational semantics (as e.g. E-APP2 uses $\llbracket T \rrbracket$). Thus, the semantic notion of execution is separated from the algorithmic notion of subtyping — and our goal in developing a complete subtyping algorithm is to ensure as many correct programs as possible are typeable. The reader may be surprised to see that FT does not include a first-class notion of function value (i.e. a term of the form $\lambda x.t$). This avoids a well-known problem of circularity in the definitions (i.e. where the semantic definition of types depends on the operational semantics and vice-versa [36, 37, 33]). In short, including function values adds unnecessary complexity and is therefore omitted. Instead, functions are declared explicitly and a runtime environment, Δ , is used to maintain the mapping from declared functions to their bodies.

An example FT program and its evaluation is given below:

$$\begin{aligned}
& f(\text{any } x) = \text{if}(x \text{ is int}) 1 \text{ else } 0 \\
& \quad \text{in } (f \ 1, f \ (1, 2)) \\
\hookrightarrow & f(\text{any } x) = \text{if}(x \text{ is int}) 1 \text{ else } 0 \\
& \quad \text{in } (\text{if}(1 \text{ is int}) 1 \text{ else } 0, f \ (1, 2)) \\
\hookrightarrow & f(\text{any } x) = \text{if}(x \text{ is int}) 1 \text{ else } 0 \\
& \quad \text{in } (1, f \ (1, 2)) \\
\hookrightarrow & f(\text{any } x) = \text{if}(x \text{ is int}) 1 \text{ else } 0 \\
& \quad \text{in } (1, \text{if}((1, 2) \text{ is int}) 1 \text{ else } 0) \\
\hookrightarrow & f(\text{any } x) = \text{if}(x \text{ is int}) 1 \text{ else } 0 \\
& \quad \text{in } (1, 0) \\
\hookrightarrow & (1, 0)
\end{aligned}$$

This example illustrates a few interesting aspects of Figure 1. Firstly, for simplicity, the order of evaluation for tuples is undefined under E-TUP. This could easily be specified, but is not important here. Secondly, the term $\text{if}(x \text{ is } T) \ t_1 \ \text{else} \ t_2$ implements a runtime type test (similar to e.g. Java's `instanceof` operator). The left-hand side of this operator is restricted to a variable, rather than a general term. This succinctly captures the problem of retyping a variable within the true (resp. false) branches of the conditional.

2.3 Flow-Typing Rules

The flow-typing rules are given in Figure 2. These are presented as judgements of the form $\Gamma \vdash t : T$, which can be read as saying: *term* t *can be shown to have type* T *under environment* Γ . The environment maps variable names to their current type, and also function names to a pair $T_1 \rightarrow T_2$ capturing the declared parameter and inferred return type. For simplicity, we assume that function names and parameter names do not intersect.

Rules T-INT, T-VAR and T-TUP are straightforward and do not warrant further discussion. The remaining rules are more interesting, and we now consider them in more detail:

- Rule T-APP. For a function application, the type of the argument is determined recursively, whilst the function's declared parameter and inferred return types are obtained from the environment. The rule checks the argument type (i.e. T_1) is a subtype of the declared parameter type (i.e. T_2) using the subtype operator (i.e. $T_1 \leq T_2$). The subtype operator will be discussed in more detail below.
- Rule T-DEC. For a function declaration, the return type is inferred by typing the body (i.e. t_2) using the current environment updated to map the parameter (i.e. x) to its declared type (i.e. T_1). Using this, the type of the outer term (i.e. t_3) is then

Flow-Typing:	
$\frac{v \in \mathbb{Z}}{\Gamma \vdash v : \text{int}}$	(T-INT)
$\frac{\Gamma(x) = T}{\Gamma \vdash x : T}$	(T-VAR)
$\frac{\Gamma \vdash t_1 : T_1, \dots, \Gamma \vdash t_n : T_n}{\Gamma \vdash (t_1, \dots, t_n) : (T_1, \dots, T_n)}$	(T-TUP)
$\frac{\Gamma \vdash t_1 : T_1 \quad \Gamma(f) = T_2 \rightarrow T_3 \quad T_1 \leq T_2}{\Gamma \vdash f t_1 : T_3}$	(T-APP)
$\frac{\Gamma[x \mapsto T_1] \vdash t_2 : T_2 \quad \Gamma[f \mapsto T_1 \rightarrow T_2] \vdash t_3 : T_3}{\Gamma \vdash f(T_1 x) = t_2 \text{ in } t_3 : T_3}$	(T-DEC)
$\frac{\Gamma[x \mapsto \Gamma(x) \wedge T_1] \vdash t_2 : T_2 \quad \Gamma[x \mapsto \Gamma(x) \wedge \neg T_1] \vdash t_3 : T_3}{\Gamma \vdash \text{if}(x \text{ is } T_1) t_2 \text{ else } t_3 : T_2 \vee T_3}$	(T-IF)

Fig. 2. Flow-typing rules for FT.

determined. Observe that, under this rule, recursive function calls cannot be typed as f is not included when typing t_2 — however, this is of little relevance to the problem being addressed.

- Rule T-IF. For a type test, the true and false branches are typed using updated environments. For the true branch, the variable being tested (i.e. x) is mapped to the intersection of its current type and that of the type test (i.e. to $\Gamma(x) \wedge T_1$) — this captures the fact that its values are known to be in both $\Gamma(x)$ and T_1 . Similarly, for the false branch, the variable being tested is mapped to the intersection of its current type and that of the negated type test (i.e. to $\Gamma(x) \wedge \neg T_1$) — this captures the fact that its values are known to be in $\Gamma(x)$ but not in T_1 . The resulting type of the type test is then the most precise type which includes the types determined for each branch (i.e. $T_2 \vee T_3$).

Observe that, in rule T-IF, the type of the tested variable may be determined as `void` for either branch — which, in such case, indicates that branch is unreachable. A modern compiler would most likely report such a situation as a syntax error (but this is an orthogonal issue).

Discussion. Having considered the flow-typing rules, we can now consider why certain constructs are included in our calculus. Firstly, function application is included since T-App requires a subtype test. Without this construct, there is no need for a subtyping algorithm such as presented in this paper. Secondly, tuple types are included because

Subtyping (incomplete):			
$\frac{}{T \leq \text{any}}$	[S-ANY1]	$\frac{}{\text{void} \leq T}$	[S-ANY2]
$\frac{}{\text{int} \leq \neg(T_1, \dots, T_n)}$	[S-INT1]	$\frac{}{(T_1, \dots, T_n) \leq \neg \text{int}}$	[S-INT2]
$\frac{\forall i. T_i \leq S_i}{(T_1, \dots, T_n) \leq (S_1, \dots, S_n)}$	[S-TUP1]	$\frac{n \neq m \vee \exists i. T_i \leq \neg S_i}{(T_1, \dots, T_n) \leq \neg (S_1, \dots, S_m)}$	[S-TUP2]
$\frac{\forall i. T_i \geq S_i}{\neg(T_1, \dots, T_n) \leq \neg(S_1, \dots, S_n)}$	[S-TUP3]		
$\frac{\forall i. T_i \leq S}{T_1 \vee \dots \vee T_n \leq S}$	[S-UNION1]	$\frac{\exists i. T \leq S_i}{T \leq S_1 \vee \dots \vee S_n}$	[S-UNION2]
$\frac{\exists i. T_i \leq S}{T_1 \wedge \dots \wedge T_n \leq S}$	[S-INTERSECT1]	$\frac{\forall i. T \leq S_i}{T \leq S_1 \wedge \dots \wedge S_n}$	[S-INTERSECT2]

Fig. 3. A sound but *incomplete* subtyping relation for the language of types defined in §2.1.

they make the subtyping problem harder (in fact, without tuples the subtyping problem for this system is fairly trivial).

2.4 Subtype Algorithm

Figure 2 employs an operation on types whose implementation is not immediately obvious — namely, determining whether one type subtypes another (i.e. $T_1 \leq T_2$). Indeed, there are many possible implementations of this operator and it is useful to consider two desirable properties:

Definition 2 (Subtype Soundness). A subtype operator, \leq , is sound if, for any types T_1 and T_2 , it holds that $T_1 \leq T_2 \implies \llbracket T_1 \rrbracket \subseteq \llbracket T_2 \rrbracket$.

Definition 3 (Subtype Completeness). A subtype operator, \leq , is complete if, for any types T_1 and T_2 , it holds that $\llbracket T_1 \rrbracket \subseteq \llbracket T_2 \rrbracket \implies T_1 \leq T_2$.

A subtype operator which exhibits both of these properties is said to be *sound* and *complete*. We know of no previous flow typing system which has this property. The most notable existing system is that of Tobin-Hochstadt and Felleisen, who developed a flow type system for Racket (formerly PLT Scheme) [23, 24]. Like the system presented here, this supports subtyping in the presence of unions and negations and was shown to be sound. However, subtyping in their system is not complete, meaning that many potentially typeable programs cannot be typed.

Example. Figure 3 provides a typical set of rules defining a subtype relation over the language of types from §2.1. These have a natural recursive implementation and are comparable to those of [23,24]. Rules S-ANY1, S-ANY2, S-INT1, S-INT2 and S-TUP1 are straightforward. We illustrate the remainder by example. Under S-TUP2, we have e.g. $(\text{int}, \text{int}) \leq \neg(\text{int}, \text{int}, \text{int})$ and $((\text{int}, \text{int}), \text{int}) \leq \neg(\text{int}, \text{int})$ whilst, similarly, we have $\neg(\text{any}, \text{any}) \leq \neg(\text{int}, \text{int})$ under S-TUP3. Under S-UNION1 we have e.g. $(\text{int}, \text{any}) \vee (\text{any}, \text{int}) \leq (\text{any}, \text{any})$ and e.g. $(\text{int}, \text{int}) \leq \text{int} \vee (\text{int}, \text{any})$ under S-UNION2. Finally, under S-INTERSECT1 we have e.g. $\text{int} \wedge (\text{int}, \text{int}) \leq \text{int}$ and e.g. $(\text{int}, \text{int}) \leq (\text{any}, \text{int}) \wedge (\text{int}, \text{any})$ under S-INTERSECT2.

The rules of Figure 3 can be shown as sound with respect to our semantic interpretation of types (i.e. Definition 1). However, they are evidently not complete. For example, neither of the following can be shown under Figure 3 (but are implied by Definition 1):

$$\text{any} \leq \text{int} \vee \neg \text{int}$$

$$(\text{int} \vee (\text{int}, \text{int}), \text{int}) \leq (\text{int}, \text{int}) \vee ((\text{int}, \text{int}), \text{int})$$

The rules of Figure 3 can be further extended to handle specific cases (such as those above). For example, we could add the following rules:

$$\frac{}{\text{any} \leq T \vee \neg T} \quad [\text{S-ANY3}]$$

$$\frac{T = (T_1, \dots, T_{k-1}, \bullet, T_{k+1}, T_n)}{T[\bullet \mapsto S_1 \vee \dots \vee S_n] \leq T[\bullet \mapsto S_1] \vee \dots \vee T[\bullet \mapsto S_n]} \quad [\text{S-TUP4}]$$

S-ANY3 allows $\text{any} \leq \text{int} \vee \neg \text{int}$ to be shown, while S-TUP4 captures distributivity across tuples, allowing $(\text{int} \vee (\text{int}, \text{int}), \text{int}) \leq (\text{int}, \text{int}) \vee ((\text{int}, \text{int}), \text{int})$. However, adding additional rules seems (in our experience) somewhat futile and just forces ever-more esoteric counter-examples. For example, using the above rules we still cannot show $\text{any} \leq (\text{int}, \text{int}) \vee (\text{int}, \text{int}, \text{int}) \vee (\neg(\text{int}, \text{int}) \wedge \neg(\text{int}, \text{int}, \text{int}))$ (which is implied by Definition 1). In essence, the issue is that the number and variety of possible equivalences between types make it very difficult to construct a set of complete rules. To address this, our approach first normalises types to eliminate many such equivalences, and to make them more manageable.

2.5 Problem Statement

We can now succinctly express the problem addressed in this paper, namely: *to develop a sound and complete subtype algorithm for the language of types defined in §2.1.* We know of no previous algorithm with this property for a comparable language of types. In §4, we present such an algorithm which, in the worst case, requires an exponential number of steps to answer a subtyping query (in the size of the types involved). This complements the work of Frisch *et al.* who provided an existence proof but did not present a practical algorithm [33]. Furthermore, determining whether a polynomial time subtyping algorithm exists for this system remains, to the best of our knowledge, an open problem.

<p>Positive Subtyping:</p> $\frac{}{T^+ \leq T^+} \quad (\text{S-REFLEX})$ $\frac{}{T^+ \leq \text{any}} \quad (\text{S-ANY})$ $\frac{\forall i \in \{1, \dots, n\}. T_i^+ \leq S_i^+}{(T_1^+ \dots, T_n^+) \leq (S_1^+, \dots, S_n^+)} \quad (\text{S-TUP})$
--

Fig. 4. Subtyping rules for positive atoms in FW.

3 Preliminaries

Before we present our algorithm for sound and complete subtyping over the language of types defined in §2.1, we first introduce the key concepts which underpin it. These then form the building blocks for our algorithmic developments in the following section.

3.1 Atoms

An important aspect of our algorithm is the definition of an *atom*. These are indivisible types which are split into the *positive* and *negative* atoms as follows:

Definition 4 (Type Atoms). Let T^* denote a type atom, defined as follows:

$$\begin{aligned} T^* &::= T^+ \mid T^- \\ T^- &::= \neg T^+ \\ T^+ &::= \text{any} \mid \text{int} \mid (T_1^+, \dots, T_n^+) \end{aligned}$$

Here, T^+ denotes a positive atom whilst T^- denotes a negative atom.

We can see from Definition 4 that a negative atom is simply a negated positive atom. Furthermore, the elements of tuple atoms are themselves positive atoms — which differs from the original definition of types, where an element could hold any possible type (including e.g. a union or intersection type). As we will see, one of the challenges we face lies in the process of converting from the general types of §2.1 into the more restricted forms used here. For example, $(\text{int} \vee (\text{int}, \text{int}), \text{any})$ can be converted into $(\text{int}, \text{any}) \vee ((\text{int}, \text{int}), \text{any})$ — which is a union of positive atoms.

The first building block we require is that of subtyping between atoms. For our purposes, this operation need only be defined for positive atoms (a fact which at first surprised us), but could be extended to negative atoms as well. Figure 4 presents the subtyping relation between positive atoms. These employ judgements of the form “ $T_1 \leq T_2$ ”, which are read simply as: *the set of values described by T_1 subtypes those of T_2* . The rules of Figure 4 are mostly straightforward. Furthermore, we can trivially obtain soundness and completeness for the subtype relation given in Figure 4:

Lemma 1. Let T_1^+ and T_2^+ be positive atoms. Then, $T_1^+ \leq T_2^+ \iff \llbracket T_1^+ \rrbracket \subseteq \llbracket T_2^+ \rrbracket$.

Proof. Straightforward by inspection of Definition 1 and Figure 4. \square

The second important building block is the observation that type atoms are *finitely indivisible*. That is, a type atom cannot be represented equivalently as a finite set of atoms which does not include itself:

Lemma 2 (Atom Indivisibility). Let T^+ and S_1^+, \dots, S_n^+ be positive type atoms where $\llbracket T^+ \rrbracket = \llbracket S_1^+ \cup \dots \cup S_n^+ \rrbracket$. Then, for some $1 \leq i \leq n$, we have $T^+ = S_i^+$.

Proof. Straightforward by inspection of Definitions 1 + 4. \square

The implications of Lemma 2 should not be overlooked. By construction, we have that $\forall i. \llbracket S_i^+ \rrbracket \subseteq \llbracket T^+ \rrbracket$ and, hence, T^+ is the unique canonical representative of the set it describes (i.e. $\llbracket T^+ \rrbracket$). Furthermore, given any S_1^+, \dots, S_n^+ where $\llbracket T^+ \rrbracket = \llbracket S_1^+ \cup \dots \cup S_n^+ \rrbracket$, we can quickly and easily identify T^+ using the subtype operator of Figure 4.

The third important building block we require is that of (positive) atom intersection. We let $T_1^+ \sqcap T_2^+$ denote the construction of a type representing the intersection of the values in T_1^+ with those of T_2^+ . Note that $T_1^+ \sqcap T_2^+$ produces either a positive atom or void (in the case of no intersection):

Definition 5 (Atom Intersection). Let T_1^+ and T_2^+ be positive atoms. Then, $T_1^+ \sqcap T_2^+$ is a positive atom or void determined as follows:

$$\begin{aligned}
T^+ \sqcap T^+ &= T^+ & (1) \\
\text{any} \sqcap T^+ &= T^+ & (2) \\
T^+ \sqcap \text{any} &= T^+ & (3) \\
\text{int} \sqcap (T_1^+, \dots, T_n^+) &= \text{void} & (4) \\
(T_1^+, \dots, T_n^+) \sqcap \text{int} &= \text{void} & (5) \\
(T_1^+, \dots, T_n^+) \sqcap (S_1^+, \dots, S_m^+) &= \text{void, if } n \neq m \text{ or } \exists i. T_i^+ \sqcap S_i^+ = \text{void} & (6) \\
&= (T_1^+ \sqcap S_1^+, \dots, T_n^+ \sqcap S_n^+), \text{ otherwise} & (7)
\end{aligned}$$

Observe that (2) + (3) and (4) + (5) are symmetric.

Definition 5 is straightforward. For example, $(\text{int}) \sqcap (\text{int}, \text{int}) = \text{void}$ as the number of fields differs (i.e. following Definition 1 where $\llbracket (\text{int}) \rrbracket \cap \llbracket (\text{int}, \text{int}) \rrbracket = \emptyset$). Also, $(\text{any}, \text{any}) \sqcap (\text{int}, \text{int}) = (\text{int}, \text{int})$ as expected. Note, $\text{int} \sqcap (\text{int}) = \text{void}$ since (T) is a tuple of arity-1 and is considered distinct from T . Finally, we can trivially obtain soundness and completeness for this operation:

Lemma 3. Let T_1^+ and T_2^+ be atoms. Then, $\llbracket T_1^+ \sqcap T_2^+ \rrbracket = \llbracket T_1^+ \rrbracket \cap \llbracket T_2^+ \rrbracket$.

Proof. Straightforward by inspection of Definition 1 and Definition 5. \square

3.2 Disjunctive Normal Form (DNF)

We now consider the procedure for converting a general type into a more classical Disjunctive Normal Form (DNF):

Definition 6 (DNF). Let $T \Longrightarrow^* T'$ denote the application of zero or more rewrite rules (defined below) to type T , producing a potentially updated type T' .

$$\neg\neg T \quad \Longrightarrow \quad T \quad (1)$$

$$\neg \bigvee_i T_i \quad \Longrightarrow \quad \bigwedge_i \neg T_i \quad (2)$$

$$\neg \bigwedge_i T_i \quad \Longrightarrow \quad \bigvee_i \neg T_i \quad (3)$$

$$(\bigvee_i S_i) \wedge \bigwedge_j T_j \quad \Longrightarrow \quad \bigvee_i (S_i \wedge \bigwedge_j T_j) \quad (4)$$

$$(\dots, \bigvee_i T_i, \dots) \quad \Longrightarrow \quad \bigvee_i (\dots, T_i, \dots) \quad (5)$$

$$(\dots, \bigwedge_i T_i, \dots) \quad \Longrightarrow \quad \bigwedge_i (\dots, T_i, \dots) \quad (6)$$

$$(\dots, \neg T, \dots) \quad \Longrightarrow \quad (\dots, \mathbf{any}, \dots) \wedge \neg(\dots, T, \dots) \quad (7)$$

$\text{DNF}(T) = T'$ denotes the computation $T \Longrightarrow^* T'$, such that no more rewrite rules apply.

Here, $\bigvee_i T_i$ (resp. $\bigwedge_i T_i$) represents a finite disjunction of the form $T_1 \vee \dots \vee T_n$ (resp. $T_1 \wedge \dots \wedge T_n$). The above rules convert a type into something similar to the classical notion of disjunctive normal form for logical expressions, with the key difference being that we must additionally factor unions, intersections and negations out of tuples. Rules 2+3 push negations inwards such that, for example, $\neg(T_1 \wedge T_2)$ rewrites to $\neg T_1 \vee \neg T_2$. Rule 4 factors unions out of intersections, such that e.g. $T_1 \wedge (T_2 \vee T_3)$ rewrites to $(T_1 \wedge T_2) \vee (T_1 \wedge T_3)$. Recall from §2.1 that $T_1 \wedge T_2$ is indistinguishable from $T_2 \wedge T_1$ and, hence, rule 4 is not restricted to rewriting only a leftmost union (as the presentation might suggest). Rules 5, 6 + 7 are responsible for factoring union and intersection types out of tuples. For example, $(\text{int} \vee (\text{int}, \text{int}), \text{any})$ rewrites by rule 4 to $(\text{int}, \text{any}) \vee ((\text{int}, \text{int}), \text{any})$. Similarly, $(\text{any} \wedge \neg \text{int}, \text{any})$ rewrites by rule 6 and then by rule 7 to give $(\text{any}, \text{any}) \wedge \neg(\text{int}, \text{any})$. Finally, we note that $\text{DNF}(T)$ may produce an exponential number of terms in the worst-case [38–40].

We now list several properties which can be trivially shown for the function $\text{DNF}(T)$, and more details can be found in [41]:

Lemma 4 (DNF Construction). Let T be a type where $\text{DNF}(T) = T'$. Then, T' has the form $\bigvee_i \bigwedge_j T_{i,j}^*$.

Lemma 5 (DNF Preservation). Let T be a type where $T \Longrightarrow T'$ by a rewrite rule from Definition 6. Then, $\llbracket T \rrbracket = \llbracket T' \rrbracket$.

Lemma 6 (DNF Termination). Let T be a type. Then, there exists a type T' for which no further rewrite rules from Definition 6 apply, such that $T \Longrightarrow^* T'$.

In considering Lemma 4, recall from Definition 4 that a type T^* represents a positive or negative atom. Thus, T^* is either a positive atom or a negated positive atom and may recursively contain only positive atoms.

4 Subtyping Algorithm

We now present our algorithm for sound and complete subtyping over the language of types defined in §2.1. We begin with an overview of the problem and our solution, and then proceed to progressively introduce the main pieces of the algorithm.

4.1 Overview

Let us reconsider the example subtyping algorithm presented in Figure 3. Recall that, whilst this algorithm can be shown sound, it is not complete. In particular, the following two rules are problematic:

$$\frac{\forall i. T_i \leq S}{T_1 \vee \dots \vee T_n \leq S} \text{ [S-UNION1]} \qquad \frac{\exists i. T \leq S_i}{T \leq S_1 \vee \dots \vee S_n} \text{ [S-UNION2]}$$

The problem is that examples of the form $T_1 \vee T_2 \leq T_3 \vee T_4$ where $\llbracket T_1 \vee T_2 \rrbracket \subseteq \llbracket T_3 \vee T_4 \rrbracket$ exist, but where neither S-UNION1 nor S-UNION2 can apply (and, hence, such examples cannot be shown under Figure 3). The following illustrates two such examples:

$$\text{int} \vee \neg \text{int} \leq (\text{int}, \text{int}) \vee \neg(\text{int}, \text{int}) \quad (1)$$

$$((\text{int}, \text{int}), \text{any}) \leq ((\text{int}, \text{int}), \text{int}) \vee \neg(\text{any}, \text{int}) \quad (2)$$

Another example is $(\text{int}, \text{any}) \wedge (\text{any}, \text{int}) \leq (\text{int}, \text{int})$ which exploits a similar problem with the S-INTERSECT1 rule.

The problem common to all these examples seems to be the number and variety of equivalences between types. To tackle these problems, we build our algorithm around the intuition that $\llbracket T_1 \rrbracket \subseteq \llbracket T_2 \rrbracket$ iff $\llbracket T_1 \rrbracket - \llbracket T_2 \rrbracket = \emptyset$. This requires an algorithm for computing the difference of two types, such that $T_1 - T_2 = \text{void}$ iff $\llbracket T_1 \rrbracket - \llbracket T_2 \rrbracket = \emptyset$. When types are represented as disjuncts of canonical conjuncts (referred to as *Canonicalised Disjunctive Normal Form* or DNF^+ for short), then computing their difference in a way that obtains the desired property is relatively easy. We proceed by first defining the notion of a *canonical conjunct* (§4.2) and then how one is constructed (§4.3). Finally, we show how a general type can be converted into DNF^+ (§4.4), and put the whole thing together illustrated with an example (§4.5).

4.2 Canonical Conjuncts

The first step in the canonicalisation process is to canonicalise intersections of the form $T_1 \wedge \dots \wedge T_n$. For example, $\text{int} \wedge \text{any}$ can be safely simplified to int . Likewise, $(\text{int}, \text{int}) \wedge \neg(\text{any}, \text{any})$ can be simplified to void , while $(\text{int}, \text{int}) \wedge \neg \text{int}$ can be simplified to (int, int) and, finally, $(\text{any}, \text{int}) \wedge \neg(\text{int}, \text{any})$ can be simplified to $(\text{any}, \text{int}) \wedge \neg(\text{int}, \text{int})$. As we will see, any intersection $\bigwedge_i T_i^*$ between atoms can be represented as a positive atom conjuncted with zero or more negative atoms, i.e. as $T_1^+ \wedge \neg T_2^+ \wedge \dots \wedge \neg T_n^+$.

Given the tools developed in §3.1 (i.e. Figure 4 and Definition 5), we can now formalise the notion of a *canonical conjunct* as follows:

Definition 7 (Canonical Conjunct). Let T^\wedge denote a canonical conjunct. Then, T^\wedge is a type of the form $T_1^+ \wedge \neg T_2^+ \wedge \dots \wedge \neg T_n^+$ where:

1. For every negation $\neg T_k^+$, we have $T_1^+ \neq T_k^+$ and $T_1^+ \geq T_k^+$.
2. For any two distinct negations $\neg T_k^+$ and $\neg T_m^+$, we have $T_k^+ \not\geq T_m^+$.

Recall the subtype relation, \geq , between positive atoms used in Definition 7 is given in Figure 4. Now, rule 1 makes sense if we recall that $T_1 \wedge \neg T_2$ can be thought of as $T_1 - T_2$; thus, in rule 1 we require that the amount “subtracted” from the positive atom by any given negative atom is strictly less than the total. For example, $(\text{int}, \text{int}) \wedge \neg(\text{any}, \text{any})$ is not permitted as this corresponds to the void (i.e. empty) type. Likewise, $(\text{any}, \text{int}) \wedge \neg(\text{int}, \text{any})$ is not permitted either since this is more precisely represented as $(\text{any}, \text{int}) \wedge \neg(\text{int}, \text{int})$. Rule 2 prohibits negative atoms from subsuming each other. For example, $(\text{any}, \text{any}) \wedge \neg(\text{int}, \text{int}) \wedge \neg(\text{any}, \text{int})$ is more precisely represented as $(\text{any}, \text{any}) \wedge \neg(\text{any}, \text{int})$. Note, we need not worry about atoms that overlap but where neither subsumes the other (i.e. where $\llbracket T_1^+ \rrbracket \cap \llbracket T_2^+ \rrbracket \neq \emptyset$ but $\llbracket T_1^+ \rrbracket \not\subseteq \llbracket T_2^+ \rrbracket$ and $\llbracket T_1^+ \rrbracket \not\supseteq \llbracket T_2^+ \rrbracket$) — this follows from Lemma 2 (indivisibility) as such types canonically represent distinct sets (hence must be retained in the conjunct).

We can make the following strong statement about canonical conjuncts based on Definition 7 — namely, that canonical conjuncts are indeed canonical:

Lemma 7 (Canonical Conjuncts). Let $T^\wedge = T_1^+ \wedge \bigwedge_i \neg T_i^+$ and $S^\wedge = S_1^+ \wedge \bigwedge_j \neg S_j^+$ be canonical conjuncts. Then, it follows that $\llbracket T^\wedge \rrbracket = \llbracket S^\wedge \rrbracket \iff T^\wedge = S^\wedge$.

Proof. Follows from proof of atom indivisibility (Lemma 2). See [41] for details. \square

4.3 Conjunct Construction

We now develop the mechanism for constructing a canonical conjunct from an arbitrary conjunct of atoms:

Definition 8 (Conjunct Canonicalisation). Let $\bigwedge_i T_i^* \implies^* \bigwedge_j S_j^*$ denote the application of zero or more rewrite rules (defined below) to $\bigwedge_i T_i^*$, producing a potentially updated version $\bigwedge_j S_j^*$.

$$\begin{array}{lll}
\text{void} \wedge \dots & \implies & \text{void} & (1) \\
T_i^+ \wedge T_j^+ \wedge \dots & \implies & (T_i^+ \sqcap T_j^+) \wedge \dots & (2) \\
T_x^+ \wedge \neg T_y^+ \wedge \dots & \implies & \text{void} & \text{if } T_x^+ \leq T_y^+ & (3) \\
& \implies & T_x^+ \wedge \dots & \text{if } T_x^+ \sqcap T_y^+ = \text{void} & (4) \\
& \implies & T_x^+ \wedge \neg(T_x^+ \sqcap T_y^+) \wedge \dots & \text{if } T_x^+ \not\geq T_y^+ & (5) \\
\neg T_x^+ \wedge \neg T_y^+ \wedge \dots & \implies & \neg T_x^+ \wedge \dots & \text{if } T_x^+ \geq T_y^+ & (6)
\end{array}$$

Let $\text{CAN}(\bigwedge_i T_i^*) = \bigwedge_j S_j^*$ denote the computation $\bigwedge_i T_i^* \implies^* \bigwedge_j S_j^*$, such that no further rewrite rules apply.

In considering the rules from Definition 8, we must recall that $T_1 \wedge T_2$ is not distinguishable from $T_2 \wedge T_1$. Therefore e.g. rule (2) picks two arbitrary positive atoms from $\bigwedge_i T_i^*$, not just the leftmost two (as the presentation might suggest). Rule 2 simply combines all the positive atoms together using the intersection operator for positive atoms (Definition 5). After repeated applications of rule 2, there will be at most one positive atom remaining. Rule 3 catches the case when the negative contribution exceeds the positive contribution (e.g. $\text{int} \wedge \neg \text{any} \implies \text{void}$). Rule 4 catches negative components which lie outside the domain (e.g. $\text{int} \wedge \neg(\text{int}, \text{int}) \implies \text{int}$). Rule 5 covers the case of negative components which lie partially outside the domain (e.g. $(\text{any}, \text{int}) \wedge \neg(\text{int}, \text{any}) \implies (\text{any}, \text{int}) \wedge \neg(\text{int}, \text{int})$). Finally, rule 6 catches the case where one negative component is completely consumed by another (e.g. $(\text{any}, \text{any}) \wedge \neg(\text{int}, \text{any}) \wedge \neg(\text{int}, \text{int}) \implies (\text{any}, \text{any}) \wedge \neg(\text{int}, \text{any})$).

Lemma 8. *Let $\bigwedge_i T_i^*$ be an arbitrary conjunct of atoms containing at least one positive atom. Then, $\text{CAN}(\bigwedge_i T_i^*)$ is either a canonical conjunct or void .*

Proof. Straightforward by case analysis on the ways in which an arbitrary conjunct of atoms does not meet the requirements of Definition 7 (and, for each case, that a rule from Definition 8 applies). See [41] for details. \square

The requirement in Lemma 8 for at least one positive atom arises because the rules of Definition 8 do not introduce positive atoms, but canonical conjuncts require them. In fact, we can easily ensure an arbitrary conjunct of atoms, $\bigwedge_i T_i^*$, has at least one positive atom — we simply add any to give $\text{any} \wedge \bigwedge_i T_i^*$.

Definition 9 (Conjunct Intersection). *Let $T_1^\wedge, \dots, T_n^\wedge$ be canonical conjuncts. Then, $T_1^\wedge \sqcap \dots \sqcap T_n^\wedge$ denotes their intersection, and is defined as $\text{CAN}(T_1^\wedge \wedge \dots \wedge T_n^\wedge)$.*

Observe that, by construction, $\prod_i T_i^\wedge$ yields either a canonical conjunct or void . Since a canonical conjunct cannot represent void , we have the required property that $\llbracket \bigwedge_i T_i^\wedge \rrbracket = \emptyset \iff \prod_i T_i^\wedge = \text{void}$. To see why a canonical conjunct cannot represent void , recall that void is short-hand for $\neg \text{any}$. Thus, we might consider $\text{any} \wedge \neg \text{any}$ to be a canonical conjunct representing void — but, this is invalid as, for a type $T_1 \wedge \neg T_2$ to be a canonical conjunct, Definition 7 requires $T_1 > T_2$. In fact, by construction, no canonical conjunct T^\wedge exists where $\llbracket T^\wedge \rrbracket = \llbracket \text{void} \rrbracket$. Finally, the following ensures the overall canonicalisation process is sound:

Lemma 9. *Let $T_1^\wedge, \dots, T_n^\wedge$ be canonical conjuncts. Then, $\llbracket \bigwedge_i T_i^\wedge \rrbracket = \llbracket \prod_i T_i^\wedge \rrbracket$.*

Proof. Straightforward by case analysis on the rules of Definition 8 to show that each rule preserves the described semantic set before and after the rewrite. See [41] for details. \square

4.4 Canonicalised Disjunctive Normal Form (DNF⁺)

Finally, we can now formally define the process for converting an arbitrary type (as defined in §2.1) into the variant of disjunctive normal form we refer to as *Canonicalised Disjunctive Normal Form* (DNF⁺):

Definition 10 (DNF⁺). Let T^\vee denote a type in Canonicalised Disjunctive Normal Form (DNF⁺). Then, either T^\vee has the form $\bigvee_i T_i^\wedge$ or is `void`.

In our definition of DNF⁺, we must include a special case for when $T = \text{void}$ since (as discussed earlier) `void` is not a canonical conjunct. We can now easily construct types in DNF⁺ as follows:

Definition 11 (DNF⁺ Construction). Let T be a type where $T' = \text{DNF}(T)$ and, hence, by Lemma 4 we have $T' = \bigvee_i \bigwedge_j T_{i,j}^*$. Then, $\text{DNF}^+(T) = \bigvee_i \prod_j T_{i,j}^*$.

In considering Definition 11, we must recall our assumption from §2.1 that $T_1 \vee T_1$ is indistinguishable from T_1 . This is important as it ensures that, if all the intersected conjuncts give `void`, then the overall result is `void` (i.e. since `void` \vee `void` = `void`, etc). This reflects our overall goal of ensuring $\llbracket T \rrbracket = \emptyset \iff \text{DNF}^+(T) = \text{void}$. We now present the overall theorem of this paper:

Theorem 1. Let T be a type (as defined in §2.1). Then, $\llbracket T \rrbracket = \llbracket \text{DNF}^+(T) \rrbracket$.

Proof. Follows immediately from Lemma 9 and Definition 11. □

4.5 Putting It All Together

We can now give a formal definition for our subtyping operator which is sound and complete for the language of types defined in §2.1, and which replaces Figure 3:

Definition 12 (Subtyping). Let T_1 and T_2 be types (as defined in §2.1). Then, it follows that $T_1 \leq T_2 \iff \text{DNF}^+(T_1 \wedge \neg T_2) = \text{void}$.

The proof that this rule is sound and complete follows immediately from Theorem 1. Furthermore it is important to realise that, whilst Definition 12 offers an improvement over Figure 3 in terms of completeness, this comes at a cost. More specifically, the subtype relation defined in Figure 3 can be implemented with a polynomial time algorithm, whilst our replacement (i.e. Definition 12) requires exponential time in the worst case. This is because the first step of the process which converts T into disjunctive normal form (i.e. Definition 6) can produce an exponential explosion in the number of terms [38–40]. As such, determining whether a sound and complete polynomial time subtyping algorithm exists remains an open problem.

We now return to consider a simple example from §2.4, and illustrate how it is resolved:

$$\begin{aligned}
 \text{any} \leq \text{int} \vee \neg \text{int} & \iff \text{DNF}^+(\text{any} \wedge \neg(\text{int} \vee \neg \text{int})) = \text{void} \\
 \text{DNF}^+(\text{any} \wedge \neg(\text{int} \vee \neg \text{int})) & = \text{CAN}(\text{DNF}(\text{any} \wedge \neg(\text{int} \vee \neg \text{int}))) \\
 & = \text{CAN}(\text{any} \wedge \neg \text{int} \wedge \text{int}) \\
 & = \text{void}
 \end{aligned}$$

Therefore, the algorithm correctly concludes that $\text{any} \leq \text{int} \vee \neg \text{int}$ holds.

5 Related Work

Tobin-Hochstadt and Felleisen consider the problem of typing previously untyped Racket (aka Scheme) programs using a flow-typing algorithm [23, 24]. Their system retypes variables within expressions dominated by type tests. However, they employ only union types and do not consider intersections or negations, making their system significantly more conservative than presented here. The work of Guha *et al.* focuses on flow-sensitive type checking for JavaScript [25]. 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. Their system employs union types, but does not support negations or intersections. Furthermore, it can retype variables as a result of runtime type tests, although only simple forms are permitted.

Since locals and stack locations are untyped in Java Bytecode, the Java Bytecode Verifier employs flow typing to ensure type safety [31]. The verifier retypes variables after assignments, but does not retype them after `instanceof` tests. And, instead of supporting explicit unions, it computes the least upper bound of the types for each variable at a meet point. A well-known problem, however, is that Java’s subtype relation does not form a complete lattice [32]. This arises because two classes can share the same super-class and implement the same interface; thus, they may not have a unique least upper bound. The solution adopted by the bytecode verifier ignores interfaces entirely and, instead, maps them to `java.lang.Object`. This approach is conservative and means some programs will fail to verify that we might otherwise expect to pass. Several works on formalising the bytecode verifier have proposed the use of intersection types as an alternative solution [42, 43].

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 [12]. 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# [15]. Here, variables are annotated with `NonNull` 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 `NonNull` 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 [13]. Pominville *et al.* also briefly discuss a flow-sensitive non-null analysis built using Soot, which does retype variables after `v!=null` checks [21]. The JACK tool for verifying `@NonNull` type annotations extends the bytecode verifier with an extra level of indirection called *type aliasing* [14]. This enables the system to retype a variable `x` as `@NonNull` in the body of an `if(x!=null)` conditional. The algorithm is formalised using a flow-sensitive type system operating on Java bytecode. JavaCOP provides an expressive language for writing type system extensions, including non-null types [22]. 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 `v` to a non-null variable if this is the first statement after a `v!=null` conditional.

Finally, there have been some attempts to incorporate intersection and union types into mainstream languages. The most relevant is that of Büchi and Weck who introduce *compound types* in to Java to overcome limitations caused by a lack of multiple inheritance [44]. Another interesting work is that of Igarashi and Nagira, who introduce union types into Java to increase code reusability [45]. Likewise, Plümicke introduces intersection types into Java to ensure methods have principle types [46]. This helps to alleviate the burden of writing complex types in some situations.

6 Conclusion

Flow-typing systems often require complex type systems involving unions, intersections and/or negations. For example, unions are often used to describe the types of variables at meet points. Likewise, intersections and negations can describe the effect of runtime type tests. However, subtype testing is a challenging algorithmic problem for a type system containing these features. In particular, to ensure the greatest number of programs as possible can be typed, we desire that subtyping is both *sound* and *complete*. Frisch *et al.* demonstrated that this problem was decidable [33]. However, their proof was not constructive and did not lend itself naturally to an implementation. In this paper, we presented a sound and complete algorithm for subtyping in the presence of unions, intersections and negations. This contrasts with previous flow type systems (e.g. [23, 24]) which are shown sound, but not complete.

We framed our algorithm in the context of a flow typing system, which is a natural fit for this work and has many well-known practical applications. Furthermore, our motivation for developing this algorithm stems from our work on the Whyley programming language [28–30], which incorporates an ambitious flow type system. However, there are other potential applications for our algorithm, such as e.g. typing XML Schema [47, 48].

Acknowledgements. The author would like to thank Sophia Drossopoulou and Nicholas Cameron for helpful comments on earlier drafts of this paper. This work is supported by the Marsden Fund, administered by the Royal Society of New Zealand.

References

1. R. Cartwright and M. Fagan. Soft typing. In *Proceedings of the ACM conference on Programming Language Design and Implementation (PLDI)*, pages 278–292, 1991.
2. D. Ancona, M. Ancona, A. Cuni, and N. D. Matsakis. RPython: a step towards reconciling dynamically and statically typed OO languages. In *Proceedings of the Dynamic Languages Symposium (DLS)*, pages 53–64, 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 *Proceedings of the ACM conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, pages 117–136, 2009.
7. J. R. Hindley. The principal type-scheme of an object in combinatory logic. *Transactions of the AMS*, 146:29–60, 1969.
8. Robin Milner. A theory of type polymorphism in programming. *Journal of Computer and System Sciences*, 17:348–375, 1978.
9. The Scala programming language. <http://lamp.epfl.ch/scala/>.
10. G. Bierman, E. Meijer, and M. Torgersen. Lost in translation: formalizing proposed extensions to C#. In *Proceedings of the ACM conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, pages 479–498, 2007.
11. D. Remy and J. Vouillon. Objective ML: An effective object-oriented extension to ML. *Theory And Practice of Object Systems*, 4(1):27–50, 1998.
12. Jeffrey S. Foster, Tachio Terauchi, and Alex Aiken. Flow-sensitive type qualifiers. In *Proceedings of the ACM conference on Programming Language Design and Implementation (PLDI)*, pages 1–12, 2002.
13. Torbjörn Ekman and Görel Hedin. Pluggable checking and inferencing of non-null types for Java. *Journal of Object Technology*, 6(9):455–475, 2007.
14. C. Male, D. J. Pearce, A. Potanin, and C. Dymnikov. Java bytecode verification for @Non-Null types. In *Proceedings of the conference on Compiler Construction (CC)*, pages 229–244, 2008.
15. M. Fähndrich and K. R. M. Leino. Declaring and checking non-null types in an object-oriented language. In *Proceedings of the ACM conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, pages 302–312, 2003.
16. Andrew C. Myers. JFlow: Practical mostly-static information flow control. In *Proceedings of the ACM symposium on the Principles Of Programming Languages (POPL)*, pages 228–241, 1999.
17. Sebastian Hunt and David Sands. On flow-sensitive security types. In *Proceedings of the ACM symposium on the Principles Of Programming Languages (POPL)*, pages 79–90, 2006.
18. A. Russo and A. Sabelfeld. Dynamic vs. static flow-sensitive security analysis. In *Proc. CSF*, pages 186–199, 2010.
19. David J. Pearce. JPure: a modular purity system for Java. In *Proceedings of the conference on Compiler Construction (CC)*, volume 6601 of *LNCS*, pages 104–123, 2011.
20. Jeffrey S. Foster, Manuel Fähndrich, and Alexander Aiken. A theory of type qualifiers. In *Proceedings of the ACM conference on Programming Language Design and Implementation (PLDI)*, pages 192–203, 1999.
21. P. Pominville, F. Qian, R. Vallée-Rai, L. Hendren, and C. Verbrugge. A framework for optimizing Java using attributes. In *Proceedings of the conference on Compiler Construction (CC)*, pages 334–554, 2001.
22. C. Andreae, J. Noble, S. Markstrum, and T. Millstein. A framework for implementing pluggable type systems. In *Proceedings of the ACM conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*, pages 57–74, 2006.
23. Sam Tobin-Hochstadt and Matthias Felleisen. The design and implementation of typed Scheme. In *Proceedings of the ACM symposium on the Principles Of Programming Languages (POPL)*, pages 395–406, 2008.
24. Sam Tobin-Hochstadt and Matthias Felleisen. Logical types for untyped languages. In *Proceedings of the ACM International Conference on Functional Programming (ICFP)*, pages 117–128, 2010.
25. A. Guha, C. Saftoiu, and S. Krishnamurthi. Typing local control and state using flow analysis. In *Proceedings of the European Symposium on Programming (ESOP)*, pages 256–275, 2011.

26. Johnni Winther. Guarded type promotion: eliminating redundant casts in Java. In *Proceedings of the Workshop on Formal Techniques for Java-like Programs*, pages 6:1–6:8, 2011.
27. What's new in Groovy 2.0? <http://www.infoq.com/articles/new-groovy-20>.
28. The Whiley programming language, <http://whiley.org>.
29. D. Pearce and J. Noble. Implementing a language with flow-sensitive and structural typing on the JVM. *Electronic Notes in Computer Science*, 279(1):47–59, 2011.
30. D. J. Pearce, N. Cameron, and J. Noble. Whiley: a language with flow-typing and updateable value semantics. Technical Report ECSTR12-09, Victoria University of Wellington, 2012.
31. Tim Lindholm and Frank Yellin. *The Java Virtual Machine Specification*. Addison Wesley, second edition, 1999.
32. X. Leroy. Java bytecode verification: algorithms and formalizations. *Journal of Automated Reasoning*, 30(3/4):235–269, 2003.
33. A. Frisch, G. Castagna, and V. Benzaken. Semantic subtyping: Dealing set-theoretically with function, union, intersection, and negation types. *Journal of the ACM*, 55(4):19:1–19:64, 2008.
34. Alexander Aiken and Edward L. Wimmers. Type inclusion constraints and type inference. In *Proc. FPCA*, pages 31–41, 1993.
35. Flemming M. Damm. Subtyping with union types, intersection types and recursive types. volume 789 of *LNCS*, pages 687–706. 1994.
36. Castagna and Frisch. A gentle introduction to semantic subtyping. In *Proceedings of the International Colloquium on Automata, Languages and Programming (ICALP)*, pages 198–199, 2005.
37. A. Frisch, G. Castagna, and V. Benzaken. Semantic subtyping. In *Proceedings of the ACM/IEEE Symposium on Logic In Computer Science (LICS)*, pages 137–146, 2002.
38. M. R. Garey and D. S. Johnson. *Computers and intractability; a guide to the theory of NP-completeness*. W.H. Freeman, 1979.
39. Christopher Umans. The minimum equivalent DNF problem and shortest implicants. *Journal of Computer and System Sciences*, 63, 2001.
40. David Buchfuhrer and Christopher Umans. The complexity of boolean formula minimization. *Journal of Computer and System Sciences*, 77(1):142–153, 2011.
41. D. J. Pearce. Sound and complete flow typing with unions, intersections and negations. Technical Report ECSTR12-20, Victoria University of Wellington, 2012.
42. Allen Goldberg. A specification of Java loading and bytecode verification. In *Proc. CCS*, pages 49–58, 1998.
43. Cornelia Pusch. Proving the soundness of a Java bytecode verifier specification in Isabelle/HOL. In *Proceedings of the conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 89–103, 1999.
44. Martin Büchi and Wolfgang Weck. Compound types for java. In *Proceedings of the ACM conference on Object-Oriented Programming, Systems, Languages and Applications (OOP-SLA)*, pages 362–373, 1998.
45. Atsushi Igarashi and Hideshi Nagira. Union types for object-oriented programming. In *Proceedings of the Symposium on Applied Computing (SAC)*, pages 1435–1441, 2006.
46. Martin Plümicke. Intersection types in java. In *Proceedings of the conference on Principles and Practices of Programming in Java (PPPJ)*, pages 181–188, New York, NY, USA, 2008. ACM.
47. Haruo Hosoya and Benjamin C. Pierce. XDuce: A statically typed XML processing language. *ACM Transactions on Internet Technology*, 3(2):117–148, 2003.
48. Véronique Benzaken, Giuseppe Castagna, and Alain Frisch. CDuce: An XML-centric general-purpose language. In *Proceedings of the ACM International Conference on Functional Programming (ICFP)*, pages 51–63, 2003.