

Delft University of Technology
Software Engineering Research Group
Technical Report Series

Language-Independent Type-Dependent Name Resolution

Hendrik van Antwerpen, Pierre Neron, Andrew Tolmach,
Eelco Visser, Guido Wachsmuth

Report TUD-SERG-2015-006

TUD-SERG-2015-006

Published, produced and distributed by:

Software Engineering Research Group
Department of Software Technology
Faculty of Electrical Engineering, Mathematics and Computer Science
Delft University of Technology
Mekelweg 4
2628 CD Delft
The Netherlands

ISSN 1872-5392

Software Engineering Research Group Technical Reports:

<http://www.se.ewi.tudelft.nl/techreports/>

For more information about the Software Engineering Research Group:

<http://www.se.ewi.tudelft.nl/>

```
@techreport{TUD-SERG-2015-006,  
  title = {Language-Independent Type-Dependent Name Resolution},  
  author = {Hendrik van Antwerpen and Pierre Neron and Andrew P. Tolmach  
    and Eelco Visser and Guido Wachsmuth},  
  year = {2015},  
  month = {July},  
  researchr = {http://researchr.org/publication/TUD-SERG-2015-006},  
  institution = {Delft University of Technology, Software Engineering Research Group},  
  number = {TUD-SERG-2015-006},  
  address = {Delft, The Netherlands},  
}
```

© copyright 2015, Software Engineering Research Group, Department of Software Technology, Faculty of Electrical Engineering, Mathematics and Computer Science, Delft University of Technology. All rights reserved. No part of this series may be reproduced in any form or by any means without prior written permission of the publisher.

Language-Independent Type-Dependent Name Resolution

Hendrik van Antwerpen

Delft University of Technology
h.vanantwerpen@student.tudelft.nl

Pierre Neron

Delft University of Technology
p.j.m.neron@tudelft.nl

Andrew Tolmach

Portland State University
apt@cs.pdx.edu

Eelco Visser

Delft University of Technology
visser@acm.org

Guido Wachsmuth

Delft University of Technology
guwac@acm.org

Abstract

We extend and combine two existing declarative formalisms, the scope graphs of Neron et al. and type constraint systems, to build a language-independent theory that can describe both name and type resolution for realistic languages with complex scope and typing rules. Unlike conventional static semantics presentations, our approach maintains a clear separation between scoping and typing concerns, while still being able to handle language constructs, such as class field access, for which name and type resolution are necessarily intertwined. We define a constraint scheme that can express both typing and name binding constraints, and give a formal notion of constraint satisfiability together with a sound algorithm for finding solutions in important special cases. We describe the details of constraint generation for a model language that illustrates many of the interesting resolution issues associated with modules, classes, and records. Our constraint generator and solver have been implemented in the Spoofox Language Workbench.

1. Introduction

Name resolution and type resolution are two fundamental concerns in programming language specification and implementation. Name resolution means determining the identifier declaration corresponding to each identifier use in a program. Type resolution means determining the type of each identifier and expression in the program, as part of performing type checking or inference. (We consider only statically-typed languages.) These two tasks are essential components of many language processing tools, including interpreters, compilers and IDEs. Moreover, precise descriptions of name and type resolution are essential parts of a formal language semantics. Yet there are as yet no universally accepted formalisms that support both specification and implementation of these tasks. This is in notable contrast to the situation with syntax definition, for which context-free grammars provide a

well-established declarative formalism that underpins a wide variety of useful tools.

In this paper, we show how two existing formalisms, scope graphs and type constraints, can be extended and combined to begin filling this gap. Our formalisms: (i) have a clear and clean underlying theory; (ii) handle a broad range of common language features; (iii) are declarative, but are realizable by practical algorithms and tools; (iv) are factored into language-specific and language-independent parts, to maximize re-use; and (v) apply to erroneous programs (for which resolution fails or is ambiguous) as well as to correct ones. Moreover, although name and type resolution are obviously related, as far as possible we treat them as separate concerns; this improves modularity and helps clarify exactly what the relationships between these two tasks are.

Our starting point is recent work by Neron *et al.* [8], which shows how name resolution for lexically-scoped languages can be formalized in a way that meets the criteria above. The name binding structure of a program is captured in a *scope graph* which records identifier declarations and references and their scoping relationships, while abstracting away program details. Its basic building blocks are *scopes*, which correspond to sets of program points that behave uniformly with respect to resolution. Each scope can contain identifier declarations and references, each tagged with its position in the original AST. Scopes can be connected by edges representing lexical nesting or import of named collections of declarations such as modules or records. A scope graph is constructed from the program AST using a language-dependent traversal, but thereafter, it can be processed in a largely language-independent way. A *resolution calculus* gives a formal definition of what it means for a reference to identifier x at position i to resolve to a declaration of x at position j , written $x_i^R \mapsto x_j^D$. Resolutions are described as paths in the scope graph obeying certain (language-specific) criteria; a given reference may resolve to one, none, or many declarations. There is a *resolution algorithm* that computes the set of declarations to which each

reference resolves, and is sound and complete with respect to the calculus.

On the type resolution side, we adopt a well-known approach based on extracting *constraints* on types and type variables from the AST and then using *unification* to solve the constraints and instantiate the variables. This technique goes back at least to Milner’s seminal paper on polymorphism [7], and has since been extended to cover many additional language features, notably subtyping. Pottier and Remy [11] give a detailed exposition, and show how an efficient resolution algorithm can be expressed using rewrite rules. The constraint approach is most commonly used for polymorphic type inference, but even for the simpler problem of monomorphic type checking, passing to constraints is a useful way to separate the language-dependent part of the task (generating the constraints) from the language-independent part (solving the constraints).

The most important connection between identifier names and static types is that each identifier should be associated with a unique type (or, in a polymorphic language, a unique type scheme). Scope graphs do not include explicit type information. However, if the language associates types with identifier declarations, it is easy to obtain the type of an identifier reference by first resolving the reference to a declaration and then looking up the associated type information by position in the AST. This simple two stage approach—name resolution using the scope graph followed by a separate type resolution stage—will work for many language constructs. But the full story is more complicated, because sometimes name resolution also depends on type resolution. Consider the program fragments in Figure 1, written in a language with nominal records and using standard dot notation for record field access. (Subscripts on identifiers represent source code positions and are not part of the language itself.) In order to resolve the type of $y_7.x_8$ we must first resolve the field name x_8 to the appropriate declaration field (x_2 or x_6). But this name resolution depends on the *type* of y_7 , so we must resolve that type first, which again, requires first resolving the *name* of y_7 . In general, we may need arbitrarily deep recursion between the two kinds of resolution. For example, to handle the nested record dereference on the last line, we must first resolve the name of y_9 , then its type, then the name and type of a_{10} , and finally the name and type of x_{11} .

To solve this challenge, we reformulate the task of generating a scope graph from a given program as one of finding a minimal solution to a set of *scope constraints* obtained by an AST traversal. Scope constraints are analogous to typing constraints, but are resolved using a different (and simpler) algorithm. We then introduce a class of *scope variables* and modify the resolution calculus to characterize resolution in potentially *incomplete* scope graphs (i.e., graphs characterized by constraints involving unresolved scope variables). We can then interleave (partial) scope graph resolution and

```

record A1{ x2: Int }
record B3{ a4: A5 x6:Bool }
...
y7.x8 // what is the type of y7 ?
y9.a10.x11 // what are the types of y9, y9.a10 ?

```

Figure 1. Program with records

type unification until a complete instantiation of all variables (types, positions, and scopes) is obtained. This approach permits us to resolve all the names and types for the record examples of Figure 1 and for a broad range of other language constructs.

Contributions Our specific contributions are as follows:

- We show how to complement name resolution based on scope graphs with type resolution based on type constraints including type-dependent name resolution (Section 2).
- We extend the name resolution calculus and algorithm of [8] to include a new kind of scope graph edge and to handle incomplete scope graphs (Section 3, Section 6).
- We define a constraint language that can express both monomorphic typing and name binding constraints, parameterized by an underlying notion of type compatibility, and define satisfiability for problems in this language (Section 4).
- We describe the details of constraint generation for a model language that illustrates many of the interesting resolution issues associated with modules, classes, and records (Section 5).
- We describe an algorithm for solving problems in our constraint language instantiated to use nominal subtyping, and show that it is sound with respect to the satisfiability specification (Section 6).

The implementation provides name and type resolution in the IDE generated with the Spoofox Language Workbench [5] for the LMR model language used in this paper and has been used to generate the scope graphs and type constraints for the examples in this paper automatically.

2. Combining Scope Graphs with Types

In this section we describe our approach to type-dependent name resolution using examples in a small model language. We show how scope graphs are used to model name binding, combine scope graphs with type constraints to model type resolution, and discuss extension of the two models to handle type-dependent name resolution.

2.1 Example Language

We illustrate the ideas using LMR (Language with Modules and Records), which extends the LM (Language with Modules) of [8]. The language does not aspire to be a real programming language, but is designed to represent typical and

<code>prog</code>	<code>= decl*</code>
<code>decl</code>	<code>= module Id { decl* } import Qid def bind</code> <code> record Id sup? { fdecl* }</code>
<code>sup</code>	<code>= extends Qid</code>
<code>fdecl</code>	<code>= id : ty</code>
<code>ty</code>	<code>= Int Qid ty → ty</code>
<code>exp</code>	<code>= int true false qid exp ⊕ exp</code> <code> if exp then exp else exp</code> <code> fun (id : ty) { exp } exp exp</code> <code> let bind* in exp letpar bind* in exp</code> <code> letrec tbind* in exp</code> <code> new Qid { fbind* } with exp do exp</code> <code> exp . id</code>
<code>Qid</code>	<code>= Id Qid . Id</code>
<code>qid</code>	<code>= id Qid . id</code>
<code>bind</code>	<code>= id = exp tbind</code>
<code>tbind</code>	<code>= id : ty = exp</code>
<code>fbind</code>	<code>= id = exp</code>

Figure 2. Syntax of LMR.

challenging name and type resolution idioms. The grammar of LMR is defined in Fig. 2. The basic features that LMR inherits from LM are:

- Modules and imports: modules can be nested and can import other modules.
- Various flavours of variable binding constructs: variable definitions (**def**), first-class functions (**fun**), and three flavors of let bindings.
- Declarations (modules, definitions, records) in the same module (scope) are mutually recursive.
- Qualified names allow access to the declaration in a module without import.

LMR extends LM with the following features:

- Static, monomorphic types: function arguments require explicit type annotations, but bindings of variables may be left for type inference to resolve.
- Declaration of nominal record types with inheritance and a corresponding subtyping relation on record types.
- Construction of (immutable) records with **new** using references to fields for initialization.
- Access to the fields of a record value using dot notation `e . f`.
- Implicit access to record fields using a Pascal-like **with** construct.

In the rest of this section we study name and type resolution for a selection of LMR constructs that explain the ideas of type-dependent name resolution using examples. Subsequent sections formalize these ideas.

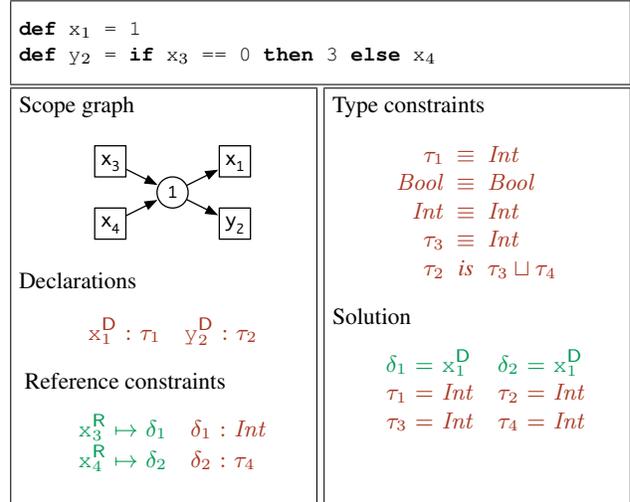


Figure 3. Declarations and references in global scope with example program, scope graph, and constraints.

2.2 Declarations and References

We recall the concepts of the scope graph approach [8] and extend it with type constraints. Consider the example in Fig. 3, which shows an LMR program (top), and the scope graph diagram and constraints (below) extracted from it. Subscripts on identifiers represent AST positions. Thus, x_1 and x_3 are different occurrences of the *same* name x .

Scope Graph The key building block of a scope graph is the *scope*, an abstraction of a set of nodes in the AST that behave uniformly with respect to name binding. In a scope graph diagram, scopes are represented by circles with numbers representing their identity. Scopes manage the visibility of *declarations*. In a diagram, declarations are represented by boxes with an *incoming arrow* from a scope. In the example program x_1 and y_2 are declarations. In constraints we denote declarations using a D superscript (e.g. x_1^D). *References* are identifiers that refer to a declaration. In diagrams, a reference is represented by means of a box with an arrow pointing to its scope. In the program x_3 and x_4 are references. In constraints we denote references with an R superscript (e.g. x_3^R). *Name resolution* in a scope graph consists of finding a path in the scope graph from a reference to a declaration. Since scope 1 contains a declaration x_1^D with the name x , both references x_3^R and x_4^R resolve to the declaration x_1^D , which we write $x_3^R \mapsto x_1^D$.

Type Constraints Scope graphs do not include explicit type information. However, by associating type information with identifier declarations, it is would be easy to obtain the type of an identifier reference by first resolving the reference to a declaration and then looking up the associated type information by position in the AST. But that would require a language-dependent mechanism, so instead we generate constraints in a generic constraint language, as illustrated in Fig. 3.

The constraints in the figure are categorized into three groups. *Declaration constraints* associate types with declarations. In the example, the constraints $x_1^D : \tau_1$ and $y_2^D : \tau_2$ associate type variables with declarations x_1^D and y_2^D . *Reference constraints* retrieve the types of variables by means of a resolution constraint associating a declaration variable to a reference, and a type association constraint for the declaration variable. For example, the constraint $x_3^R \mapsto \delta_1$ requires that reference x_3^R resolve to declaration variable δ_1 , and the constraint $\delta_1 : Int$ requires the type of that declaration to be Int because of the use of the reference in the equality operator. Finally, *type constraints* pose equality and subtype constraints on the types assigned to declarations and expressions. For example, the constraint $\tau_1 \equiv Int$ arises from the declaration of x_1^D , the constraint $Bool \equiv Bool$ arises from the condition of the `if`, the constraint $Int \equiv Int$ arises from the integer 3. (We will leave the trivial equality constraints out in further examples.) Finally, the branches of the `if` generate a least upper-bound constraint $\tau_2 \text{ is } \tau_3 \sqcup \tau_4$ on the types of the branches.

It is also useful to categorize constraints by whether they affect **name resolution** or **resolution**. To help visualize this distinction, we use two different colors; later in the paper, we add additional colors for further kinds of constraints. (But you won't lose essential information by reading this paper in grayscale, since the categorization is strictly syntactic.)

Resolution The combination of a scope graph and type constraints define a *resolution problem*. A solution for such a problem is a substitution for the declaration and type variables in the problem such that (1) name resolutions are consistent with the scope graph according to the rules of the resolution calculus (Section 3), and (2) all type constraints are satisfied. For the example, in the solution for Fig. 3, the substitution for δ_1 is dictated by the fact that the only path through the scope path starting from x_3^R ends at x_1^D , and the substitution for τ_2 is deduced from the equality constraints on τ_1 and τ_3 , the unification of τ_1 and τ_4 (via x_1^D), and the lower-bound constraint on τ_2 .

2.3 Lexical Scope and Subtypes

Fig. 4 shows a larger LMR example that illustrates lexical scope and subtype constraints.

Lexical scope is modeled using parent arrows between scopes in the scope graph. In the example, scope 3, corresponding to the body of the `fun`, is enclosed in scope 2, corresponding to the `letrec`, which is enclosed in scope 1, the global scope of the program. Resolution of a reference proceeds from the scope of the reference to parent scopes until a matching declaration is found. Thus, reference n_5^R resolves to declaration n_3^D , which shadows n_1^D .

A function application such as `fac6(n7 - 1)` requires that the type of the actual parameter (τ_7) is a subtype of the type of the formal parameter (τ_6).

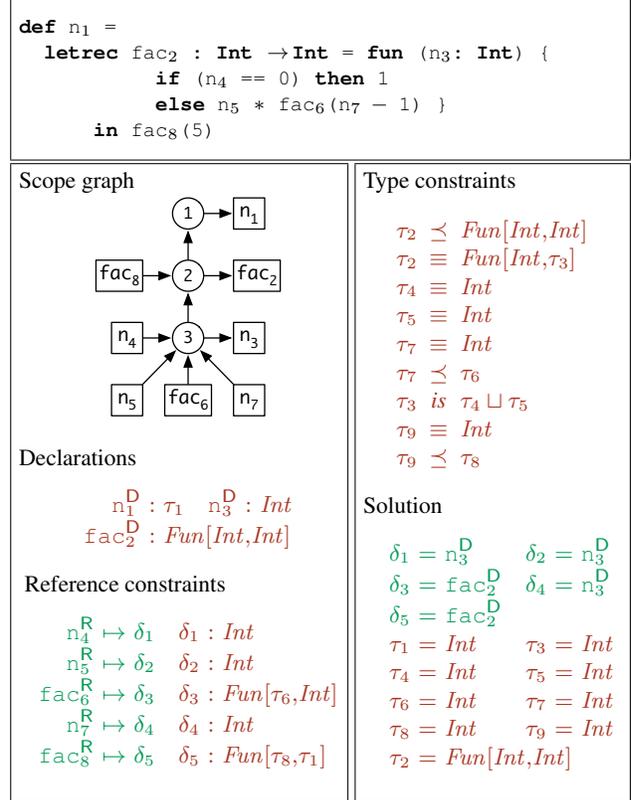


Figure 4. Lexical scoping modeled in a scope graph and subtyping relations captured in constraints.

2.4 Imports

In addition to lexical scope, many programming languages provide features for making declarations in scopes selectively available ‘at a distance’. Examples of such constructs are modules with imports in ML and classes with inheritance in Java. To model such features, scope graphs provide *associated scopes* and *imports*.

Associated Scope The LMR program in the left of Fig. 5 consists of two *modules* A_1 and B_3 and an import of the former into the latter. The declarations in these modules are contained in scopes 2 and 3, which are child scopes of the root scope 1. These scopes are *associated* with the declaration of the name of the module, which is represented in a scope graph diagram with an open arrow from the declaration (e.g. A_1^D) to the scope (e.g. 2).

Imports The declarations in a scope are only visible to references in lexically enclosed scopes, i.e. following parent edges to child scopes. An *import* makes the declarations in a scope visible in another, not necessarily lexically related, scope. An import is represented by (1) a regular reference of the name in its enclosing scope, and (2) an import in that scope. The latter is represented using an open arrow from a scope to a reference. For example, `import A4` is represented

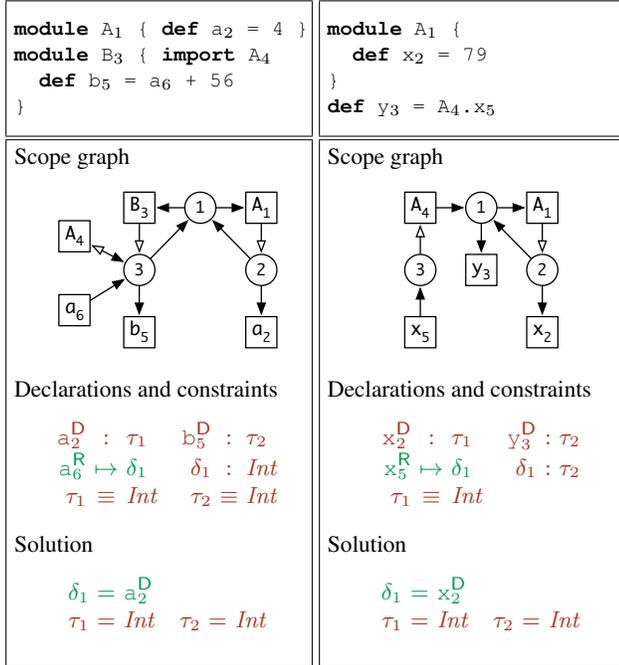


Figure 5. Module imports and qualified names with example programs, scope graphs, and constraints.

by the reference A_4^R in scope 3 and an import arrow from scope 3 to A_4^R .

Resolving through Imports Name resolution in the presence of associated scopes and imports proceeds as follows. If a scope S_1 contains an import x_i^R , which resolves to a declaration x_j^D with associated scope S_2 , then all declarations in S_2 are reachable in S_1 . Thus, in the example, reference a_6^R resolves to declaration a_2^D since the import A_4^R resolves to declaration A_1^D , and the associated scope 2 of A_1^D contains declaration a_2^D . Note that the resolution calculus is parameterized by the policy used to disambiguate conflicting resolutions. Here we use the default policy of [8] that prefers imported declarations over declarations in parents.

Qualified Names Another common pattern for accessing the declarations in a scope is through qualified names. Instead of importing *all* declarations in a scope, a single declaration is accessed. For example, in the right program from Fig. 5 the expression $A_4.x_5$ refers to the declaration x_2^D in module A_1 . This pattern can be modeled using the scope graph ingredients that we have seen so far. The reference x_5^R is defined as a reference of parentless scope 3. The only declarations visible in scope 3 are through the import of A_4^R , which is itself a reference in scope 1. Thus, since A_4^R resolves to A_1^D , the declarations in its associated scope 2 are visible in scope 3, and therefore, x_5^R resolves to x_2^D .

2.5 Type-Dependent Name Resolution

To summarize, scope graphs provide a language-independent model for formalizing the binding rules in programming

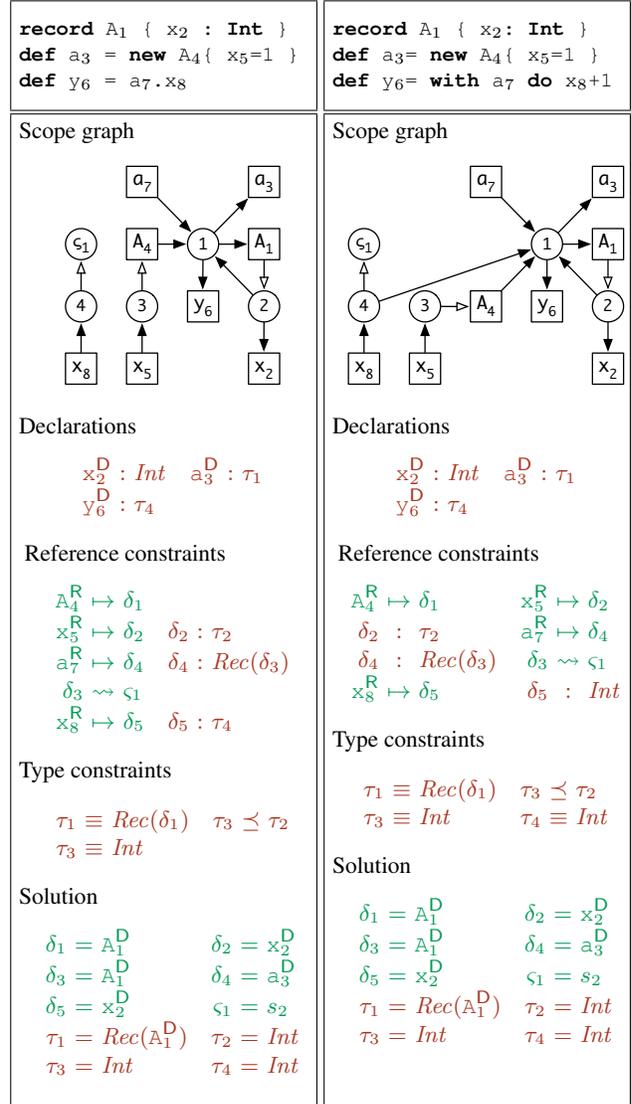


Figure 6. Field Access and **with** expression modeled by virtual scopes, reference, association, and type constraints.

languages. Neron et al. [8] show that the approach covers a wide range of name binding idioms. In this section we have shown that scope graphs can be complemented with type constraints to express the static typing requirements on programs (to be formalized later in this paper). These constraints use name resolution constraints to express the dependence of type resolution on name resolution.

However, for some language constructs the resolution of a name to its declaration depends on the type of another expression. For example, in field access expression $e.f$, in order to resolve the field f , one first needs to find the type of the expression e and then to look for f in the scope associated with the type. This scheme induces a dependency, not only of the name resolution but also of the scope graph construction (one does not know in which scope the reference f

lies) on the type resolution. We model such *type-dependent name resolution* by means of constraints over the edges in the scope graph.

Field Access Both examples in Fig. 6 illustrate the approach. In the left example, we are particularly interested in the field access in the definition of y_6^D . The reference x_8^R is a field access in the record value of a_7^R . Thus, x_8^R should be resolved in the associated scope of the type of the receiver expression a_7^R . This is similar to the resolution of a qualified name, which we modeled by resolving the qualified name in a parentless scope into which we imported the module. Thus, we create a parentless scope (4) and add x_8^R as reference in that scope. However, in this case we do not know (the name of) the record type that should be imported into the parentless scope. Therefore, we proceed as follows. We create a new scope identified by a *scope variable* ς_1 that acts as a placeholder for the scope that we want to import into the parentless scope 4. We add a *direct import edge* (open arrow) from scope 4 to scope ς_1 . Then, we resolve a_7^R using $a_7^R \mapsto \delta_4$ and obtain the type of its definition through $\delta_4 : \text{Rec}(\delta_3)$, which should be a record type identifying the record definition δ_3 . We use a new form of scope graph constraint, $\delta_3 \rightsquigarrow \varsigma_1$, to specify that ς_1 is the associated scope of δ_3 . Solving these constraints will lead to a solution for ς_1 — in this case the associated scope of a_1^D , scope 2 — such that the appropriate scope can be imported into scope 4. After that x_8^R can be resolved as usual to its definition $x_8^R \mapsto \delta_5$, which leads to its type $\delta_5 : \tau_4$.

Note that scope 3 and related edges and constraints model the resolution of the field initializer in the definition of a_3^D , which is similar to the pattern for qualified names, but applies to a list of initializer expressions.

With As final example, we discuss an expression form inspired by the **with** statement in the Pascal language. In the expression **with** e **do** e' , the fields of the record value of e are added to the lexical environment of e' . That is, a variable reference x in e' will be interpreted as a field of the record value when the record has indeed a field with name x ; otherwise the variable is considered as a regular reference in the enclosing lexical context. Static resolution again requires resolving variables in e' in the associated scope of the record type of e , but this time defaulting to the enclosing lexical scope.

Fig. 6 shows on its right how this is modeled for the expression **with** a_7 **do** $x_8 + 1$ using a scope (4) that directly imports a placeholder scope (ς_1) as the lexical context for the references in the body of the **with**. The scope variable is resolved through the constraints $a_7^R \mapsto \delta_4$, $\delta_4 : \text{Rec}(\delta_3)$, and $\delta_3 \rightsquigarrow \varsigma_1$ to the associated scope of the type of a_7^R . Unlike in the case of field access the scope for the body of the **with** does have a parent scope (1), so that references that are not to fields of the record are resolved in the lexical context.

2.6 Roadmap

The rest of this paper formalizes the approach to type-dependent name resolution sketched in this section. Section 3 reviews the resolution calculus of Neron et al. [8] and extends it with direct imports between scopes. Section 4 defines the syntax and semantics of a constraint language that can be used by language front-ends to express the name binding and type rules of a language. In Section 5 we give a complete account of extraction of constraints for all LMR constructs. Section 6 describes a resolution algorithm that finds solutions for resolution problems.

3. Extended Scope Graphs

In this section we recall the formal theory of name resolution of Neron et al. [8] consisting of a scope graph model and resolution calculus, and extend the model with direct imports to model type-dependent name resolution as introduced in the previous section.

3.1 Scope Graphs

A *scope graph* is a language-independent model for representing the name binding structure of programs. A scope graph \mathcal{G} is built around three basic types of nodes derived from the program abstract syntax tree (AST), *declarations*, *references*, and *scopes*:

- A *declaration* is an occurrence of an identifier that introduces a name. x_i^D denotes the definition of name x at position i in the program. We omit the position i when this is unimportant or can be inferred from context. $\mathcal{D}(\mathcal{G})$ denotes the set of declarations of \mathcal{G} .
- A *reference* is an occurrence of an identifier referring to a declaration. We write x_i^R for a reference with name x at position i (again, we sometimes omit the position.) $\mathcal{R}(\mathcal{G})$ denotes the set of references of \mathcal{G} .
- A *scope* is an abstraction over a set of nodes in the AST that behave uniformly with respect to name binding. $\mathcal{S}(\mathcal{G})$ denotes the set of scopes of \mathcal{G} .

Given these sets, a scope graph is defined by the following functions:

- Each declaration d in $\mathcal{D}(\mathcal{G})$ is declared within a scope denoted $\mathcal{S}c(d)$.
- Each declaration d has an optional *associated scope*, $\mathcal{D}Sc(d)$ that is the scope corresponding to the body of the declaration. For example, the declarations in a module are elements of its associated scope.
- Each reference r in $\mathcal{R}(\mathcal{G})$ is declared within a scope denoted $\mathcal{S}c(r)$.
- Each scope S in $\mathcal{S}(\mathcal{G})$ has an optional *parent scope* $\mathcal{P}(S)$ that corresponds to its *lexically enclosing scope*. The parent relation has to be well-founded, i.e. there is no infinite sequence S_n such that $S_{n+1} = \mathcal{P}(S_n)$.

Resolution paths	
s	$:= \mathbf{D}(x_i^D) \mid \mathbf{I}(x_i^R, x_j^D) \mid \mathbf{I}(S) \mid \mathbf{P}$
p	$:= [] \mid s \mid p \cdot p$ (inductively generated)
$[] \cdot p$	$= p \cdot [] = p$
$(p_1 \cdot p_2) \cdot p_3$	$= p_1 \cdot (p_2 \cdot p_3)$
Example well-formed paths predicate	
	$WF(p) \Leftrightarrow p \in \mathbf{P}^* \cdot \mathbf{I}(_)^*$
Example specificity ordering on paths	
$\overline{\mathbf{D}(_) < \mathbf{I}(_)}$	(DI) $\frac{s_1 < s_2}{s_1 \cdot p_1 < s_2 \cdot p_2}$ (Lex1)
$\overline{\mathbf{I}(_) < \mathbf{P}}$	(IP) $\frac{p_1 < p_2}{s \cdot p_1 < s \cdot p_2}$ (Lex2)
$\overline{\mathbf{D}(_) < \mathbf{P}}$	(DP)

Figure 7. Resolution paths, well-formedness predicate, and specificity ordering as introduced in [8]

Edges in scope graph	
$\frac{\mathcal{P}(S_1) = S_2}{\mathbb{I} \vdash \mathbf{P} : S_1 \rightarrow S_2}$	(P)
$\frac{y_i^R \in \mathcal{I}(S_1) \setminus \mathbb{I} \quad \mathbb{I} \vdash p : y_i^R \mapsto y_j^D}{\mathbb{I} \vdash \mathbf{I}(y_i^R, y_j^D) : S_1 \rightarrow \mathcal{DSc}(y_j^D)}$	(I)
$\frac{S_2 \in \mathcal{IS}(S_1)}{\mathbb{I} \vdash \mathbf{I}(S_2) : S_1 \rightarrow S_2}$	(D)
Transitive closure	
$\overline{\mathbb{I} \vdash [] : A \rightarrow A}$	(N)
$\frac{\mathbb{I} \vdash s : A \rightarrow B \quad \mathbb{I} \vdash p : B \rightarrow C}{\mathbb{I} \vdash s \cdot p : A \rightarrow C}$	(T)
Reachable declarations	
$\frac{x_i^D \in \mathcal{D}(S') \quad \mathbb{I} \vdash p : S \rightarrow S' \quad WF(p)}{\mathbb{I} \vdash p \cdot \mathbf{D}(x_i^D) : S \rightarrow x_i^D}$	(R)
Visible declarations	
$\frac{\mathbb{I} \vdash p : S \rightarrow x_i^D \quad \forall j, p' (\mathbb{I} \vdash p' : S \rightarrow x_j^D \Rightarrow \neg(p' < p))}{\mathbb{I} \vdash p : S \mapsto x_i^D}$	(V)
Reference resolution	
$\frac{\mathcal{Sc}(x_i^R) = S \quad \{x_i^R\} \cup \mathbb{I} \vdash p : S \mapsto x_j^D}{\mathbb{I} \vdash p : x_i^R \mapsto x_j^D}$	(X)

Figure 8. Resolution calculus from [8] extended with direct import rule D

- Each scope S has an associated set of references $\mathcal{I}(S)$, that represents the *imports* in this scope

We define by comprehension the set of declarations enclosed in a scope S , as $\mathcal{D}(S) = \{d \mid \mathcal{Sc}(d) = S\}$. We occasionally subscript these defining functions by the particular graph \mathcal{G} to which they apply, writing e.g. $\mathcal{DSc}_{\mathcal{G}}$.

3.2 Resolution Calculus

Given this model, the *resolution calculus* defines the *resolution* of a reference to a declaration in a scope graph [8] as the minimal path from reference to declaration through parent and import edges. A path p is a list of steps representing the atomic scope transitions in the graph. A step is a parent step \mathbf{P} , an import step $\mathbf{I}(y^R, y^D)$ where y^R is the import used and y^D its corresponding declaration, or a declaration step $\mathbf{D}(x^D)$ leading to a declaration x^D . Given a seen import set \mathbb{I} , a path p is a valid resolution in the graph from reference x_i^R to declaration x_j^D when the following statement holds:

$$\mathbb{I} \vdash p : x_i^R \mapsto x_j^D$$

The calculus in Fig. 8 defines the resolution relation in terms of edges in the scope graph, reachable declarations, and visible declarations.

The resolution calculus is parameterized by two predicates on paths, a *well-formedness predicate* $WF(p)$ and an *ordering relation* $<$ that allows the formalization of different name-binding policies such as transitive vs non-transitive imports. A typical definition of the well-formedness predicate is *no-parents-after-imports*, which entails that a resolution cannot proceed to a lexical parent after an import transition. Fig. 7 presents the definition of paths (p) consisting of steps (s) and examples of a path well-formedness predicate and a path ordering relation. This configuration supports arbitrary levels of lexical scope (\mathbf{P}^*), transitive imports ($\mathbf{I}(_)^*$), no-parents-after-imports (an $\mathbf{I}(_)$ step cannot be followed by a \mathbf{P}), prefer local declarations over imported declarations (DI), prefer local declarations over declarations in parents (DP), and prefer imported declarations over declarations in parents (IP).

3.3 Direct Imports

In order to model type-dependent name resolution we extend the scope graph with *direct imports*. A direct import defines a direct link between two scopes without the use of a reference. In addition to its set of associated imports (references of the form x^R), a scope is extended with an associated set of directly imported scopes $\mathcal{IS}(S)$. For these imports we introduce the (D) scope transition rule, which is similar to the (I) rule of the original calculus, except that this transition does not require the intermediate resolution of a reference:

$$\frac{S_2 \in \mathcal{IS}(S_1)}{\mathbb{I} \vdash \mathbf{I}(S_2) : S_1 \rightarrow S_2} \quad (D)$$

The complete resolution calculus with this new rule is presented in Figure 8.

4. Constraint Language

In this section we introduce the syntax and declarative semantics of constraints.

4.1 Syntax of Constraints

Fig. 9 defines the syntax of constraints. The language independent base terms of the constraint language are:

- *Declarations* in D , which are either ground declarations x_i^D of the program or variables δ
- *References* in R , which are ground references x_i^R of the program
- *Scopes* in S , which are either ground scopes of the program denoted S_i or variables ς
- *Types* in T , which are either type variables τ or type constructor applications $c(T, \dots, T)$ with $c \in C_{\mathcal{T}}$, the set of language-specific type constructors.

Given these terms we define the syntax of constraints, which come in two flavors, *facts* and proper *constraints*. *Facts*, defined by the sort F , correspond to known facts about a program: the scope of a reference ($Sc(R) := S$), the scope of a declaration ($Sc(D) := S$), the associated scope of a declaration ($D \rightsquigarrow S$), the parent of a scope ($P(S) := S$), a named import in a scope ($S \in \mathcal{I}(S)$), a direct import in a scope ($S \in \mathcal{IS}(S)$), and a subtype relation between types ($T <: T$). Proper *constraints*, defined by the sort C , represent the restrictions on name and type resolution, which consist of: resolution of a reference to a declaration ($R \mapsto D$), equality of two types ($T \equiv T$), subtype relation between two types ($T \preceq T$), associated scope of a declaration ($D \rightsquigarrow S$), the type of a declaration ($D : T$), and the least upper bound of two sorts ($T \text{ is } T \sqcup T$). Facts and constraints can be combined using conjunction ($C \wedge C$) and True represents the trivially satisfiable constraint. As before, we use different colors to help distinguish between **facts about the scope graph**, **facts about the subtyping relation**, **proper typing constraints**, and **proper resolution constraints**.

Language-Specific Types The language of constraints defined above is independent of the language under analysis, except for the type constructors introduced by this language. Therefore we assume a set of language specific types constructors $C_{\mathcal{T}}$ and each constructors c has an associated arity $c :: n$. For example, Int and Bool are type constructors with arity 0 and Fun is a type constructor with arity 2. To represent user-defined types, such as classes in object-oriented languages or algebraic data types in functional languages, a type constructor can also include the scope-graph declaration corresponding to the type definition. For example, record types in LMR are represented by $\text{Rec}(d)$ with d a declaration in the program. Thus, in Fig. 6, the record definition A_1 defines the type $\text{Rec}(A_1^D)$

$A :=$	$Sc(R) := S$	$D \rightsquigarrow S$	$S \in \mathcal{I}(S)$
	$Sc(D) := S$	$P(S) := S^?$	$S \in \mathcal{IS}(S)$
	$T <: T$		
$C :=$	True	$R \mapsto D$	$T \equiv T$
	$C \wedge C$	$D \rightsquigarrow S$	$T \preceq T$
	$D : T$	$T \text{ is } T \sqcup T$	A
$D :=$	$\delta \mid x_i^D$	$S :=$	$\varsigma \mid S_i \quad S^? := S \mid \perp$
$R :=$	x_i^R	$T :=$	$\tau \mid c(T, \dots, T)$ with $c \in C_{\mathcal{T}}$

Figure 9. Syntax of constraints

4.2 Semantics of Constraints

In our approach, the abstract syntax tree of a program p is reduced by a language-specific extraction function $\llbracket p \rrbracket$ to a constraint following the syntax defined in Fig. 9. Given commutativity and associativity of the conjunction operator, such a constraint is equivalent to one of the form

$$F_1 \wedge \dots \wedge F_n \wedge C_1 \wedge \dots \wedge C_m$$

consisting of a set of facts F_i and a set of proper constraints C_j . (We define an example extraction function in the next section.) The facts define the scope graph and subtyping relation with respect to which the proper constraints need to be solved.

Interpretation of Facts We denote by $F^{<:}$ the set of facts of the form $T_1 <: T_2$ in F that will be used to build the corresponding subtyping relation. We denote by $F^{\mathcal{G}}$ the subset formed by the other facts that will define the scope graph of the program. Given a ground set of facts F , we denote by $|F|$ the interpretation of F as the pair \mathcal{G}, \leq where \leq is the subtyping relation derived from $F^{<:}$ and \mathcal{G} is the scope graph derived from $F^{\mathcal{G}}$.

Subtyping From the set of facts $F^{<:}$ we derive the relation \leq between ground types, built using the type constructors in $C_{\mathcal{T}}$. We require each arguments in the signature of a constructor c to be annotated with a variance parameter. Thus the signature of a type constructor c is declared as $c :: v_1 * \dots * v_n$, where the v_i are variance annotations. A variance v is a non-empty subset of $\{-, +\}$ written $-$ for contravariant, $+$ for covariant and \pm for invariant. We also denote \leq by \leq^+ , \geq by \leq^- and $=$ by \leq^{\pm} . Given such signatures for all the type constructors, we define the subtyping relation \leq derived from a set $F^{<:}$ of subtyping facts by the following inductive rules:

$$\frac{}{T \leq T} \quad \frac{T_1 \leq T_2 \quad T_2 \leq T_3}{T_1 \leq T_3}$$

$$\frac{T_1 <: T_2 \in F^{<:}}{T_1 \leq T_2} \quad \frac{c :: v_1 * \dots * v_n \quad \forall i, s_i \leq^{v_i} t_i}{c(s_1, \dots, s_n) \leq c(t_1, \dots, t_n)}$$

Scope graph The set of facts $F^{\mathcal{G}}$ define the scope graph of the program. These facts define the set of scopes, dec-

larations and references and the corresponding relations as follows:

- $P(S) := S'$ defines a new scope S and declares its parent $\mathcal{P}(S)$ as S' when S' is not \perp
- $\mathcal{S}c(x^D) := S$ defines a new declaration x^D and declares its enclosing scope $\mathcal{S}c(x^D)$ as S
- $d \rightsquigarrow S$ declares scope S as the associated scope $\mathcal{D}\mathcal{S}c(d)$ of declaration d
- $\mathcal{S}c(x^R) := S$ defines a new reference x^R and declares its enclosing scope $\mathcal{S}c(x^R)$ as S
- $r \in \mathcal{I}(S)$ adds the reference r to the set of named imports $\mathcal{I}(S)$ of scope S
- $S' \in \mathcal{I}\mathcal{S}(S)$ adds the scope S' to the set of direct imports $\mathcal{I}(S)$ of scope S

The result is a correct scope graph according to Section 3 provided that the parent relation is well-founded.

Interpretation of Proper Constraints The interpretation of a proper constraints (which we will just call constraints) is defined as a truth value in a context, which is a triple of the following elements:

- A scope graph \mathcal{G} , as defined in Section 3
- A subtyping relation \leq on ground types
- A typing environment ψ mapping declarations in $\mathcal{D}(\mathcal{G})$ to types in \mathbb{T}

A context \mathcal{G}, \leq, ψ satisfies a constraint C if the predicate $\mathcal{G}, \leq, \psi \models C$ holds. This predicate is defined by the set of inductive rules in Fig. 10, where $=$ is the syntactic equality on terms, $\vdash_{\mathcal{G}} x_i^R \mapsto x_j^D$ is the resolution relation for graph \mathcal{G} and, when it exists, $\sqcup_{\leq} S$ denotes the least upper bound of types in S according to order \leq .

4.3 Program Resolution

The goal of the resolution of the program p is to build a multi-sorted substitution ϕ and a typing environment ψ such that, if $\llbracket p \rrbracket = F_p \wedge C_p$ then the following property holds:

$$|\phi(F_p)|, \psi \models \phi(C_p) \quad (\diamond)$$

Where $\phi(E)$ denotes the application of the substitution ϕ to all the variables appearing in E that are in the domain of ϕ . Note that ϕ has to make $\phi(F_p)$ a set of ground facts in order to be able to interpret it whereas some free variables may remain in $\phi(C_p)$. When the proposition \diamond holds we say that ψ and ϕ resolve p .

5. Constraint Collection

In this section, we show how to collect constraints for name resolution and typechecking from programs in the LMR language, whose concrete syntax was given in Fig. 2. The full collection algorithm is shown in Figures 11 and 12. Collection is performed by a single traversal over the program

$\overline{\mathcal{G}, \leq, \psi \models \text{True}}$	(C-TRUE)
$\frac{\mathcal{G}, \leq, \psi \models C_1 \quad \mathcal{G}, \leq, \psi \models C_2}{\mathcal{G}, \leq, \psi \models C_1 \wedge C_2}$	(C-AND)
$\frac{\psi(d) = T}{\mathcal{G}, \leq, \psi \models d : T}$	(C-TYPEOF)
$\frac{\vdash_{\mathcal{G}} p : x_i^R \mapsto x_j^D}{\mathcal{G}, \leq, \psi \models x_i^R \mapsto x_j^D}$	(C-RESOLVE)
$\frac{\mathcal{D}\mathcal{S}c_{\mathcal{G}}(d) = S}{\mathcal{G}, \leq, \psi \models d \rightsquigarrow S}$	(C-SCOPEOF)
$\frac{t_1 = t_2}{\mathcal{G}, \leq, \psi \models t_1 \equiv t_2}$	(C-EQ)
$\frac{T_1 \leq T_2}{\mathcal{G}, \leq, \psi \models T_1 \preceq T_2}$	(C-SUBTYPE)
$\frac{T = \sqcup_{\leq} \{T_1, T_2\}}{\mathcal{G}, \leq, \psi \models T \text{ is } T_1 \sqcup T_2}$	(C-LUB)

Figure 10. Interpretation of proper constraints

that collects scope and subtyping facts, name resolution constraints, and typing constraints all in one pass. (The color codings should help in distinguishing these different kinds of constraints.)

To simplify and compress the presentation, we describe the algorithm as operating over LMR's concrete syntax. (Our actual implementation operates over the abstract syntax of LMR, and is written in DynSem, a declarative domain-specific language for expressing semantics; although readable, it is relatively verbose.) The algorithm is defined by a family of functions indexed by syntactic category (*decl*, *exp*, etc.). Each function takes a syntactic item and possibly one or more auxiliary parameters, and (usually) returns a constraint, possibly involving one or more fresh variables or new scope identifiers. Functions are defined by a set of rules, one for each possible syntactic form in the category. For example, $\llbracket - \rrbracket_s^{decl}$ has four rules (for **module**, **import**, **def** and **record** declarations, respectively), and is parameterized by the scope s into which declared identifiers are to be installed; it returns the conjunction of constraints that enforces correct name and type resolution for the declaration, some of which are derived by invoking generation functions on syntactic sub-components.

To further streamline the presentation, we use the notation $\llbracket - \rrbracket^{c^*}$ on sequences of items of syntactic category c to mean the result of applying $\llbracket - \rrbracket^c$ to each item and returning the conjunction of the resulting constraints, or True for the empty sequence. Similarly, $\llbracket - \rrbracket^{c?}$ works on an optional c item; it applies $\llbracket - \rrbracket^c$ to the item if it is present and returns True otherwise. Throughout, we use metavariable x_i for a (lower case) term variable at position i and X_i for an (upper case) module or record name at position i , with one excep-

$\llbracket n \rrbracket_{s,t}^{exp}$	$:= t \equiv Int$	
$\llbracket true \rrbracket_{s,t}^{exp}$	$:= t \equiv Bool$	$\llbracket false \rrbracket_{s,t}^{exp} := t \equiv Bool$
$\llbracket e_1 \oplus e_2 \rrbracket_{s,t}^{exp}$	$:= t \equiv t_3 \wedge \llbracket e_1 \rrbracket_{s,t_1}^{exp} \wedge \llbracket e_2 \rrbracket_{s,t_2}^{exp}$	$(\oplus \text{ has type } t_1 \times t_2 \rightarrow t_3)$
$\llbracket \text{if } e_1 \text{ then } e_2 \text{ else } e_3 \rrbracket_{s,t}^{exp}$	$:= t \text{ is } \tau_2 \sqcup \tau_3 \wedge \llbracket e_1 \rrbracket_{s,Bool}^{exp} \wedge \llbracket e_2 \rrbracket_{s,\tau_2}^{exp} \wedge \llbracket e_3 \rrbracket_{s,\tau_3}^{exp}$	$(\text{fresh } \tau_2, \tau_3)$
$\llbracket Xs.x_i \rrbracket_{s,t}^{exp}$	$:= x_i^R \mapsto \delta \wedge \delta : t \wedge \llbracket Xs.x_i \rrbracket_s^{qid}$	$(\text{fresh } \delta)$
$\llbracket e_1 e_2 \rrbracket_{s,t}^{exp}$	$:= \tau_2 \preceq \tau_1 \wedge \llbracket e_1 \rrbracket_{s,Fun[\tau_1,t]}^{exp} \wedge \llbracket e_2 \rrbracket_{s,\tau_2}^{exp}$	$(\text{fresh } \tau_1, \tau_2)$
$\llbracket e.x_i \rrbracket_{s,t}^{exp}$	$:= P(S') := \perp \wedge \varsigma \in \mathcal{IS}(S') \wedge \mathcal{SC}(x_i^R) := S' \wedge \delta_1 \rightsquigarrow \varsigma \wedge x_i^R \mapsto \delta_2$ $\wedge \delta_2 : t \wedge \llbracket e \rrbracket_{s,Rec(\delta_1)}^{exp}$	$(\text{new } S')(\text{fresh } \delta_1, \delta_2, \varsigma)$
$\llbracket \text{with } e_1 \text{ do } e_2 \rrbracket_{s,t}^{exp}$	$:= P(S') := s \wedge \varsigma \in \mathcal{IS}(S') \wedge \delta \rightsquigarrow \varsigma \wedge \llbracket e_1 \rrbracket_{s,Rec(\delta)}^{exp} \wedge \llbracket e_2 \rrbracket_{S',t}^{exp}$	$(\text{new } S')(\text{fresh } \delta, \varsigma)$
$\llbracket \text{new } Xs.X_i \{bs\} \rrbracket_{s,t}^{exp}$	$:= P(S') := s \wedge X_i^R \in \mathcal{I}(S') \wedge X_i^R \mapsto \delta \wedge t \equiv Rec(\delta)$ $\wedge \llbracket Xs.X_i \rrbracket_s^{qid} \wedge \llbracket bs \rrbracket_{s,S'}^{fbind^*}$	$(\text{new } S')(\text{fresh } \delta)$
$\llbracket x_i = e \rrbracket_{s,s'}^{fbind}$	$:= \mathcal{SC}(x_i^R) := s' \wedge x_i^R \mapsto \delta \wedge \delta : \tau_1 \wedge \tau_2 \preceq \tau_1 \wedge \llbracket e \rrbracket_{s,\tau_2}^{exp}$	$(\text{fresh } \delta, \tau_1, \tau_2)$

Figure 12. Constraint generation for LMR.

tion: for compactness, we give just one rule for both *qid* and *Qid*, in which x_i can be either kind of identifier. We write Xs for a dot-separated sequence of module or record names, which can be empty (in which case, by convention, $Xs.x$ doesn't have a leading dot).

Let us trace how the constraint generator works on some of the different syntactic forms of LMR. A complete program is a sequence of mutually-recursive top-level declarations, so $\llbracket - \rrbracket^{prog}$ creates a new root scope S in which they are to be installed, generates a fact that S is a parentless scope, and then conjoins the constraints for each declaration, passing S as a parameter to the declaration generator for *decl*. $\llbracket - \rrbracket_s^{decl}$ generates a fact that installs any declared identifier into scope s ; the rest of its behavior depends on the kind of declaration:

A **module** builds a new lexical child scope for its declarations, so the generator function creates a new scope S' and generates facts that s is the parent of S' and that the module name declaration is associated with S' . It then conjoins the constraints obtained by recursively invoking $\llbracket - \rrbracket_{S'}^{decl}$ on its member declarations.

An **import** doesn't declare an identifier, but instead generates constraints forcing the imported module name into the import set and reference set for s . The $\llbracket - \rrbracket_s^{qid}$ invocation generates additional constraints needed to describe references to potentially qualified names. (We omit discussion of the details, which are slightly complex.)

A **def** invokes an auxiliary generating function $\llbracket - \rrbracket_{s_r,s_d}^{bind}$ to process the definition; s_d is the scope into which the defined identifier's declaration should be installed, and s_r is the scope into which any identifier references in the defining expression should go. (In this invocation, the two scope parameters are the same, but the *bind* generator is invoked at other places, e.g. for the **letpar** expression, where they are not.) For bindings without explicit type annotation, a constraint

is generated giving the defined identifier a fresh type τ , and function $\llbracket - \rrbracket_{s,t}^{exp}$ is invoked to generate constraints for the defining expression with the expected type $t = \tau$. (We discuss generation of expression constraints below.) For bindings with an explicit concrete type annotation, we also generate a constraint that τ be a subtype of the declared type, after it is translated into an internal type constructor using auxiliary function $\llbracket - \rrbracket_s^{ty}$. This function, unlike all the others, returns a pair of things: the internalized type constructor, and any constraints generated by references to record names (which are installed into the scope given by parameter s).

A **record** is handled similarly to a **module**, but the details are more complicated. If the record has a super-type (non-empty **extends** clause), an auxiliary function $\llbracket - \rrbracket_{s_r,s_d,t}^{sup}$ is invoked to generate constraints to describe the inheritance. Parameter s_r is the scope (here s) in which the super-type name is to be resolved, s_d is the scope (here S') into which the super-type's scope is to be imported, and t is the new record type (here $Rec(X_i^D)$). The *sup* generator builds a subtyping fact constraint (the only source of such facts in LMR) that relates the new record type to the result of resolving the super-type name, via a fresh declaration variable δ . The fields of the new record are declared and given type constraints by another auxiliary function.

Constraint generation for expressions is largely straightforward. The *expr* generator is parameterized by the scope s into which references should be installed and the expected type t of the expression (often a type variable). Expressions introducing local bindings re-use the *bind* generator. (There is function for sequential **let**, which is desugared into nested **letpar** expressions before constraint generation is performed.) Note that generated constraints always force an expression to have a precise type, which is designed to be minimal in the subtyping hierarchy. Subtyping is allowed only at function applications, at bindings to ex-

$\llbracket ds \rrbracket^{prog}$	$:= P(S) := \perp$	$(new S)$
	$\wedge \llbracket ds \rrbracket_S^{decl*}$	
$\llbracket \text{module } X_i \{ ds \} \rrbracket_s^{decl}$	$:= Sc(X_i^D) := s$	$(new S')$
	$\wedge P(S') := s$	
	$\wedge X_i^D \rightsquigarrow S' \wedge \llbracket ds \rrbracket_{S'}^{decl*}$	
$\llbracket \text{import } X_s . X_i \rrbracket_s^{decl}$	$:= X_i^R \in \mathcal{I}(s) \wedge \llbracket X_s . X_i \rrbracket_s^{qid}$	
$\llbracket \text{def } b \rrbracket_s^{decl}$	$:= \llbracket b \rrbracket_{s,s}^{bind}$	
$\llbracket x_i = e \rrbracket_{s_r, s_d}^{bind}$	$:= Sc(x_i^D) := s_d$	$(fresh \tau)$
	$\wedge x_i^D : \tau \wedge \llbracket e \rrbracket_{s_r, \tau}^{exp}$	
$\llbracket x_i : t = e \rrbracket_{s_r, s_d}^{bind}$	$:= Sc(x_i^D) := s_d$	$(fresh \tau)$
	$\wedge x_i^D : t' \wedge \tau \preceq t'$	
	$\wedge C \wedge \llbracket e \rrbracket_{s_r, \tau}^{exp}$	
	where $\llbracket t \rrbracket_{s_r}^{ty} = (t', C)$	
$\llbracket \text{record } X_i u \{ fs \} \rrbracket_s^{decl}$	$:= Sc(X_i^D) := s$	$(new S')$
	$\wedge P(S') := s$	
	$\wedge X_i^D \rightsquigarrow S'$	
	$\wedge \llbracket u \rrbracket_{s, S', Rec(X_i^D)}^{sup?} \wedge \llbracket fs \rrbracket_{s, S'}^{fdecl*}$	
$\llbracket \text{extends } X_s . X_i \rrbracket_{s_r, s_d, t}^{sup}$	$:= X_i^R \in \mathcal{I}(s_d)$	$(fresh \delta)$
	$\wedge X_i^R \mapsto \delta$	
	$\wedge t <: Rec(\delta) \wedge \llbracket X_s . X_i \rrbracket_{s_r}^{qid}$	
$\llbracket x_i : t \rrbracket_{s_r, s_d}^{fdecl}$	$:= Sc(x_i^D) := s_d$	
	$\wedge x_i^D : t' \wedge C$	
	where $\llbracket t \rrbracket_{s_r}^{ty} = (t', C)$	
$\llbracket x_i \rrbracket_s^{qid}$	$:= Sc(x_i^R) := s$	
$\llbracket X_s . X_j . x_i \rrbracket_s^{qid}$	$:= Sc(x_i^R) := S'$	$(new S')$
	$\wedge P(S') := \perp$	
	$\wedge X_j^R \in \mathcal{I}(S') \wedge \llbracket X_s . X_j \rrbracket_s^{qid}$	
$\llbracket \text{fun } f(x_i : t) \{ e \} \rrbracket_{s, t}^{exp}$	$:= P(S') := s$	$(new S')$
	$\wedge Sc(x_i^D) := S'$	
	$\wedge x_i^D : t'$	
	$\wedge t \equiv Fun[t', \tau_2]$	$(fresh \tau_2)$
	$\wedge C \wedge \llbracket e \rrbracket_{S', \tau_2}^{exp}$	
	where $\llbracket t \rrbracket_s^{ty} = (t', C)$	
$\llbracket \text{letrec } bs \text{ in } e \rrbracket_{s, t}^{exp}$	$:= P(S') := s$	$(new S')$
	$\wedge \llbracket bs \rrbracket_{S', S'}^{bind*} \wedge \llbracket e \rrbracket_{S', t}^{exp}$	
$\llbracket \text{letpar } bs \text{ in } e \rrbracket_{s, t}^{exp}$	$:= P(S') := s$	$(new S')$
	$\wedge \llbracket bs \rrbracket_{S', S'}^{bind*} \wedge \llbracket e \rrbracket_{S', t}^{exp}$	
$\llbracket \text{Int} \rrbracket_s^{ty}$	$:= (Int, True)$	
$\llbracket \text{Bool} \rrbracket_s^{ty}$	$:= (Bool, True)$	
$\llbracket t_1 \rightarrow t_2 \rrbracket_s^{ty}$	$:= (Fun[t_1', t_2'], C_1 \wedge C_2)$	
	where $\llbracket t_1 \rrbracket_s^{ty} = (t_1', C_1)$	
	and $\llbracket t_2 \rrbracket_s^{ty} = (t_2', C_2)$	
$\llbracket X_s . X_i \rrbracket_s^{ty}$	$:= (Rec(\delta),$	$(fresh \delta)$
	$X_i^R \mapsto \delta \wedge \llbracket X_s . X_i \rrbracket_s^{qid}$	

Figure 11. Constraint generation for LMR.

explicitly annotated identifiers, and in the conditional expressions, for which a least-upper-bound constraint is generated. Scope variables ς are introduced only for field dereference and **with** expressions.

6. Resolution Algorithm

In this section, we describe an algorithm for computing program resolutions in the sense of Section 4.3. Suppose we have a program p from which we collect a set of constraints $\llbracket p \rrbracket = F_p \wedge C_p$, where F_p is a conjunction of facts and C_p is a conjunction of proper constraints. Then recall that a *resolution* for p is a multi-sorted substitution ϕ and a typing environment ψ such that

$$|\phi(F_p)|, \psi \models \phi(C_p) \quad (\diamond)$$

Our algorithm works only for a restricted class of generated constraints: all facts must be ground, except that (i) scope variables ς can appear in direct import facts (e.g. $\varsigma \in \mathcal{IS}(S)$), and (ii) type variables τ and declaration variables δ can appear on the right-hand side of a subtyping fact (e.g. $Rec(A_i^D) <: Rec(\delta)$). This restriction is met by the constraints generated by the LMR collection algorithm in Section 5. Broader classes of constraints might be useful for other languages; we defer exploration of algorithms that could handle these to future work.

6.1 Handling Variables in Facts

The basic approach of the algorithm is to apply the definitions in Section 4.2 to the facts to build a scope graph and a subtyping relation, and then use these to resolve proper constraints of the form $x^R \mapsto d$ or $t_1 \preceq t_2$ in the context of a conventional unification-based algorithm. However, since the facts can contain variables, we cannot fully define the scope graph or subtyping relation before starting constraint resolution, because we don't fully know ϕ . Thus, our algorithm builds ϕ (and Ψ) incrementally. The key idea is that we can resolve some proper constraints even when ϕ is not yet fully defined, in such a way that the resolution remains valid as it becomes more defined.

Subtyping The construction of the subtyping relation from a set of ground facts given in Section 4.2 is monotonic. Let \leq_F be the subtyping order generated from a set of ground facts F . Then given two sets of ground facts F_1 and F_2 , we have the following property:

$$F_1 \subseteq F_2 \Rightarrow T_1 \leq_{F_1} T_2 \Rightarrow T_1 \leq_{F_2} T_2$$

If F is any set of (not necessarily ground) facts, and \bar{F} is its subset of ground facts, then for all substitutions ϕ mapping type variable to ground types we have:

$$T_1 \leq_{\bar{F}} T_2 \Rightarrow T_1 \leq_{\phi(F)} T_2$$

Therefore, if we can deduce a subtyping relation between two types by only using the ground facts then this relation will still hold under any subsequent substitution.

$$\begin{aligned}
R[\mathbb{I}](x^R) &:= R_V[\{x^R\} \cup \mathbb{I}, \{\}](x, \mathcal{S}c(x^R)) \\
R_V[\mathbb{I}, \mathbb{S}](x, S) &:= R_L[\mathbb{I}, \mathbb{S}](x, S) \triangleleft R_P[\mathbb{I}, \mathbb{S}](x, S) \\
R_L[\mathbb{I}, \mathbb{S}](x, S) &:= R_D[\mathbb{I}, \mathbb{S}](x, S) \triangleleft R_I[\mathbb{I}, \mathbb{S}](x, S) \\
R_D[\mathbb{I}, \mathbb{S}](x, S) &:= \begin{cases} \emptyset & \text{if } S \in \mathbb{S} \\ \{x_i^D \mid x_i^D \in \mathcal{D}(S)\} & \end{cases} \\
R_I[\mathbb{I}, \mathbb{S}](x, S) &:= \begin{cases} \emptyset & \text{if } S \in \mathbb{S} \\ \text{exception} & \text{if } \mathcal{I}\mathcal{S}(S) \text{ contains a variable} \\ \bigcup \{R_L[\mathbb{I}, \{S\} \cup \mathbb{S}](x, S') \mid S' \in \mathcal{I}\mathcal{S}[\mathbb{I}](S) \cup \mathcal{I}\mathcal{S}(S)\} & \end{cases} \\
\mathcal{I}\mathcal{S}[\mathbb{I}](S) &:= \{\mathcal{D}\mathcal{S}c(y^D) \mid y^R \in \mathcal{I}(S) \setminus \mathbb{I} \wedge y^D \in R[\mathbb{I}](y^R)\} \\
R_P[\mathbb{I}, \mathbb{S}](x, S) &:= \begin{cases} \emptyset & \text{if } S \in \mathbb{S} \\ R_V[\mathbb{I}, \{S\} \cup \mathbb{S}](x, \mathcal{P}(S)) & \end{cases}
\end{aligned}$$

Figure 13. Name resolution algorithm

Scope Graphs The situation is a bit more complicated with respect to scope graphs. The non-strictly positive premise of the (V) rule of the resolution calculus makes the derivation of a resolution relation from a graph non-monotonic with respect to additions to the graph. For example, suppose that in some graph \mathcal{G} a reference x^R in a scope S resolves to declaration x_i^D in the parent scope S' . In a bigger graph \mathcal{G}' that also has a declaration $x_{i'}^D$ in S itself, x^R will resolve to $x_{i'}^D$, and the old resolution to x_i^D will be shadowed. Therefore we can not simply resolve a reference in a graph built from ground facts and expect this resolution to remain valid later in the resolution process.

However, we have restricted the set of constraints we handle so that almost all facts used for scope graph construction are in fact ground from the beginning. The only exception is for direct import declarations, where the imported scope can be a scope variable; recall that this construction is essential for expressing record field access, where the resolution of the field name depends on the type of the record expression. In order to handle these unknown direct imports, we define an extension of the scope graph structure, called an *incomplete scope graph*, that also allows scope variables as direct imports in addition to ground scopes. The construction of the incomplete scope graph from a set of facts with variable direct imports is similar to the one for ordinary scope graphs given in Section 4.2.

The resolution calculus as presented in Fig. 8 is only defined on ground scope graphs. Given an incomplete scope graph \mathcal{G} , a reference x^R is said to resolve to a declaration x_i^D if and only if this resolution is valid in all ground instances of this incomplete graph:

$$\vdash_{\mathcal{G}} x^R \mapsto x_i^D \triangleq \forall \sigma, \vdash_{\mathcal{G}, \sigma} x^R \mapsto x_i^D \quad (\blacklozenge)$$

where we write $\vdash_{\mathcal{G}}$ for the resolution relation for graph \mathcal{G} and \mathcal{G}, σ is the ground scope graph corresponding to the application of substitution σ to variables in \mathcal{G} . In order to be able to detect eventual duplicate resolutions in the program we also want to ensure that an incomplete graph provides *all* the possible resolutions of a given reference. In particular, if a resolution is unique in an incomplete graph, we want it to be unique in all its ground instances. An incomplete graph \mathcal{G} is stable for a reference x^R , denoted $\mathcal{G} \uparrow x^R$, if all the

resolutions in all its ground instances are the same:

$$\mathcal{G} \uparrow x^R \triangleq \forall \sigma, \sigma' \vdash_{\mathcal{G}, \sigma} x^R \mapsto x_i^D \Rightarrow \vdash_{\mathcal{G}, \sigma'} x^R \mapsto x_i^D$$

The resolution algorithm in Fig. 13 defines resolution in (potentially) incomplete scope graphs. The \triangleleft operator is defined by $S_1 \triangleleft S_2 \triangleq \text{if } S_1 \neq \emptyset \text{ then } S_1 \text{ else } S_2$. This algorithm raises an exception if the graph is not stable for the reference.

The algorithm is correct, i.e. we have

$$x_i^D \in R_{\mathcal{G}}(x^R) \implies \vdash_{\mathcal{G}} x^R \mapsto x_i^D \wedge \mathcal{G} \uparrow x^R \quad (\star)$$

where $R_{\mathcal{G}}(x^R)$ denotes the top-level resolution function $R[\emptyset](x^R)$ for the graph \mathcal{G} . We now sketch a proof of this fact. First, notice that the algorithm terminates using the lexicographic ordering $(\#(\mathcal{R}(\mathcal{G}) \setminus \mathbb{I}), \#(\mathcal{S}(G) \setminus \mathbb{S}))$, where $\#(A)$ denotes the cardinality of set A . We next prove that on ground scope graphs, this algorithm behaves like the standard resolution algorithm presented in [8]. If \mathcal{G} is ground then:

$$R_{\mathcal{G}}[\mathbb{I}](x^R) = \{x_i^D \mid \mathbb{I} \vdash_{\mathcal{G}} x^R \mapsto x_i^D\} \quad (i)$$

Proof. In this case, since the graph is ground, no exceptions can be thrown. Therefore the proof is an adaptation of Theorem 1 of [8]. The only differences in this version of the algorithm are: (a) it has an extra case for direct imports, which can be simply handled in the proof by adapting the named import case of the original proof; and (b) it computes resolution only for a single name argument rather than complete sets of visible and reachable declarations, which induces a weaker proof obligation. \square

Now let \mathcal{G} be an incomplete scope graph and \mathcal{G}' one of its instances. If a resolution on \mathcal{G} terminates with a set of declarations then the resolution on \mathcal{G}' does too:

$$R_{\mathcal{G}}[\mathbb{I}](x^R) = S \implies R_{\mathcal{G}'}[\mathbb{I}](x^R) = S \quad (ii)$$

Proof. By induction on the termination order of the algorithm $(\#(\mathcal{R}(\mathcal{G}) \setminus \mathbb{I}), \#(\mathcal{S}(G) \setminus \mathbb{S}))$. Since exceptions are never caught, and since an exception is triggered as soon as a scope variable is encountered, if a run of the algorithm on \mathcal{G} starting from x^R does terminate with a result then this run is exactly the same on \mathcal{G}' . \square

Finally, we can prove \star :

Proof. Let $S = R_{\mathcal{G}}(x^R)$ and pick $x_i^D \in S$.

To prove that x^R resolves to x_i^D in \mathcal{G} , let \mathcal{G}' be an arbitrary ground instance of \mathcal{G} . Using (ii) we have $x_i^D \in R_{\mathcal{G}'}(x^R)$ and by (i) we have $\vdash_{\mathcal{G}'} x^R \mapsto x_i^D$. By \blacklozenge , we get that $\vdash_{\mathcal{G}} x^R \mapsto x_i^D$.

To prove stability, let \mathcal{G}_1 and \mathcal{G}_2 be ground instances of \mathcal{G} . Then by (ii), $R_{\mathcal{G}_1}(x^R) = S = R_{\mathcal{G}_2}(x^R)$, so by definition we have $\mathcal{G} \uparrow x^R$. \square

6.2 Constraint Solving Algorithm

In Fig. 14 we present an algorithm to solve the constraint system from Section 4. The algorithm is a non-deterministic rewrite system working over tuples $(C, \mathcal{G}, F^{<:}, \psi)$ of a constraint, a scope graph, a set of subtyping facts, and a typing environment. It is non-deterministic in the sense that rules may be applied to any atomic constraint in any order considering that \wedge is associative and commutative.

Name resolution introduces ambiguity, since a reference x^R may resolve to multiple definitions. If this happens the solver branches, picking a different resolution for x^R in every branch. The returned solution is a set of all the $(C, \mathcal{G}, F^{<:}, \psi)$ tuples the solver was able to construct. The initial state of the solver is the collected constraint, the (incomplete) scope graph built from the scope graph facts, the subtyping facts, and an empty typing environment. The algorithm will eliminate clauses from C while instantiating \mathcal{G} and $F^{<:}$ and filling ψ . The algorithm terminates when the constraint is empty or no more clauses can be solved. Each rule solves one constraint, possibly updating components of the tuple or applying a substitution to it.

The S-RESOLVE rule solves $x^R \mapsto \delta$ constraints using the resolution algorithm from Fig. 13. If a resolution is found, it is substituted for the variable δ . If the scope graph is incomplete, the algorithm might throw an exception, in which case the constraint is left to be solved later.

The S-ASSOC rule solves $x^D \rightsquigarrow \varsigma$ constraints, by looking up the scope S associated with ground declaration x^D in the scope graph. By substituting S for ς , the scope graph becomes more complete, possibly allowing more references to be resolved.

Rule S-EQUAL solves equality constraints $T_1 \equiv T_2$. It uses first order unification $\mathcal{U}(T_1, T_2)$, as described in [1]. The resulting substitution is applied to the tuple.

Rule S-SUBTYPE solves constraints of the form $t_1 \preceq t_2$ by checking that $t_1 \preceq_{F^{<:}} t_2$ for the ground types t_1 and t_2 . The check might not succeed if $F^{<:}$ still contains variables, in which case it might be solved later.

Rule S-LUB solves T is $t_1 \sqcup t_2$ constraints. It does so by calculating the least upper bound $t = (t_1 \sqcup t_2)$ of the ground types t_1 and t_2 and generating a new equality constraint $T \equiv t$. The solver depends here on a language-specific least upper bound function \sqcup , which for LMR is presented in Fig. 15 in the Appendix.

Constraints of the form $x^D : T$ are solved by rule S-TYPEOF. The first rule is used the first time x^D is encountered and just adds it to the typing environment. For every next encounter, the other rule unifies the type T from the constraint with the type $\psi(x^D)$ from the typing environment.

The trivial constraint True is handled by S-TRUE.

6.3 Correctness

We want to prove the soundness of the constraint resolution algorithm, that is, that the solver produces a correct solution to the program resolution problem. If the solver reduces to an empty set of constraints, then the initial constraint was satisfiable. Moreover we want to ensure that the produced typing environment is a valid one, that is, it corresponds to a solution. Therefore we want to ensure the following property:

$$\begin{aligned} & \forall C, \mathcal{G}, F^{<:}, \psi, \mathcal{G}', F^{<:'}, \psi', \\ & (C, \mathcal{G}, F^{<:}, \psi) \longrightarrow^* (\text{True}, \mathcal{G}', F^{<:'}, \psi') \Rightarrow \\ & \exists \sigma, \sigma(\mathcal{G}), \leq_{\sigma(F^{<:})}, \psi' \models \sigma(C_1) \quad (\diamond) \end{aligned}$$

Proof. To prove this result we first state some results on the auxiliary unification and least upper bound computations.

Unification If $\mathcal{U}(t_1, t_2) = \sigma$ then $\sigma t_1 = \sigma t_2 \wedge \sigma \sigma = \sigma$. See [1] for a survey on unification problem and unification algorithms for first order terms.

Least Upper Bound Similarly, given a set of ground subtyping facts F , if $(t_1 \sqcup t_2) \xrightarrow{F} t$ then t is the least upper bound of t_1 and t_2 for \leq_F , i.e. $t = \sqcup_{\leq_F} \{t_1, t_2\}$. For LMR, the least upper bound computation is presented in Fig. 15 in the Appendix.

Resolution Soundness We now can prove the property \diamond of the constraint resolution algorithm. We first prove that for each reduction step, if the output is satisfiable the input is also satisfiable in the same definition-to-type environment. This is stated by the following property:

$$\begin{aligned} & \forall (C_1, \mathcal{G}_1, F^{<:}_1, \psi_1), (C_2, \mathcal{G}_2, F^{<:}_2, \psi_2), \\ & (C_1, \mathcal{G}_1, F^{<:}_1, \psi_1) \longrightarrow (C_2, \mathcal{G}_2, F^{<:}_2, \psi_2) \Rightarrow \\ & \forall \sigma, (\sigma \mathcal{G}_2, \leq_{\sigma(F^{<:}_2)}, \sigma \psi_2) \models \sigma(C_2) \Rightarrow \\ & \exists \sigma', (\sigma' \mathcal{G}_1, \leq_{\sigma'(F^{<:}_1)}, \sigma' \psi_2) \models \sigma'(C_1) \quad (\dagger) \end{aligned}$$

The proof of this property is by case analysis on the reduction step and is presented in Appendix A.1.

Using this result \dagger , we can prove property \diamond by a simple induction on the number of reduction steps. \square

7. Related Work

There are several ideas and efforts that deal directly or indirectly with the interaction between typing and name binding. These efforts are usually in the context of a specific language or formalism. We have not found a *language-independent*

$(x^R \mapsto \delta \wedge C, \mathcal{G}, F^{<}, \psi) \longrightarrow [\delta \mapsto x^D](C, \mathcal{G}, F^{<}, \psi)$	(S-RESOLVE)
<i>where $x^D \in R_G(x^R)$ without exception</i>	
$(x^D \rightsquigarrow_{\zeta} \wedge C, \mathcal{G}, F^{<}, \psi) \longrightarrow [\zeta \mapsto S](C, \mathcal{G}, F^{<}, \psi)$	(S-ASSOC)
<i>where $\mathcal{D}Sc_G(x^D) = S$</i>	
$(T_1 \equiv T_2 \wedge C, \mathcal{G}, F^{<}, \psi) \longrightarrow \sigma(C, \mathcal{G}, F^{<}, \psi)$	(S-EQUAL)
<i>where $\mathcal{U}(T_1, T_2) \longrightarrow \sigma$</i>	
$(T \text{ is } t_1 \sqcup t_2 \wedge C, \mathcal{G}, F^{<}, \psi) \longrightarrow (T \equiv t \wedge C, \mathcal{G}, F^{<}, \psi)$	(S-LUB)
<i>where $F^{<}$ is ground and $(t_1 \sqcup t_2) \xrightarrow{F^{<}} t$</i>	
$(t_1 \preceq t_2 \wedge C, \mathcal{G}, F^{<}, \psi) \longrightarrow (C, \mathcal{G}, F^{<}, \psi)$	(S-SUBTYPE)
<i>where $t_1 \leq_{F^{<}} t_2$</i>	
$(x^D : T \wedge C, \mathcal{G}, F^{<}, \psi) \longrightarrow \begin{cases} (C, \mathcal{G}, F^{<}, \{x^D \mapsto T\} \cup \psi) & \text{if } x^D \notin \text{dom}(\psi) \\ (\psi(x^D) \equiv T \wedge C, \mathcal{G}, F^{<}, \psi) & \text{else} \end{cases}$	(S-TYPEOF)
$(\text{True} \wedge C, \mathcal{G}, F^{<}, \psi) \longrightarrow (C, \mathcal{G}, F^{<}, \psi)$	(S-TRUE)

Figure 14. Constraint solving algorithm

approach to formalizing the interaction. A proposal to add type-directed name resolution [10] to Haskell identifies the dependency between type inference and name resolution as a possible problem. Introduction of a name-resolution constraint in the type checker to defer name resolution is mentioned as a possible solution. In Java, member names are resolved based on nominal types. In formal treatments for Java-like languages such as Jinja [6] and Featherweight Java [4], this is done by building a type-members mapping and using a lookup function in the typing rules. In our approach a custom mechanism is unnecessary; member resolution is just a special case of a uniform approach to handling name resolution. The JastAdd Java compiler [3] uses reference attribute grammars to express the name analysis of Java programs. While the attribute definitions provide clean design patterns for complex name binding problems, they do not provide reusable language-independent abstractions. Indeed, the patterns for tree traversal for name look-up in JastAddJ, provided some of the inspiration for the scope graph and resolution calculus abstraction.

Type Inference Algorithms The origin of type inference using constraints and the corresponding algorithm W goes back to Damas and Milner in [7, 2]. Wand simplified it in [17] and it has then been extended to support more complex type systems including records [12], constrained types to handle subtyping [15], GADTs [14, 13] and type classes [16]. The HM(X) system [9] is a generalization of the Hindley/Milner system parameterized in the constraint domain X, it is thoroughly described by Pottier and Remy in [11]. However, all of these constraint systems are often presented in an extension of the lambda calculus with relatively simplistic name binding constructs. Our current presentation does

not support any kind of generalization over type variables. In future work we would like to lift our connection between types and name binding to handle more complex type systems such as the ones listed above, combining the power of name resolution using scope graphs and the expressivity of these type systems.

8. Conclusion

We have presented a theory that combines extended scope graphs with type constraints to support language-independent specification of the name binding and typing concerns of programming languages. We have implemented a proof of concept constraint generator and solver, and used it as analysis framework in the Spoofox Language Workbench, applying it to LMR, a model language with interesting interactions between name binding and typing.

Further research directions include proving completeness of the constraint resolution algorithm (on suitably restricted sets of constraints); extending the theory with operators to express additional requirements on solutions, such as uniqueness of declarations; and applying the approach with more advanced type-system features, such as parametric polymorphism.

References

- [1] F. Baader and T. Nipkow. *Term rewriting and all that*. Cambridge University Press, 1998.
- [2] L. Damas and R. Milner. Principal type-schemes for functional programs. In *POPL*, pages 207–212, 1982.
- [3] T. Ekman and G. Hedin. Modular name analysis for java using jastadd. In *GTTSE*, pages 422–436, 2006.

- [4] A. Igarashi, B. C. Pierce, and P. Wadler. Featherweight java: a minimal core calculus for java and gj. *TOPLAS*, 23(3):396–450, 2001.
- [5] L. C. L. Kats and E. Visser. The Spoofox language workbench: rules for declarative specification of languages and IDEs. In *OOPSLA*, pages 444–463, 2010.
- [6] G. Klein and T. Nipkow. A machine-checked model for a java-like language, virtual machine, and compiler. *TOPLAS*, 28(4):619–695, 2006.
- [7] R. Milner. A theory of type polymorphism in programming. *jcss*, 17(3):348–375, 1978.
- [8] P. Neron, A. P. Tolmach, E. Visser, and G. Wachsmuth. A theory of name resolution. In *ESOP*, April 2015.
- [9] M. Odersky, M. Sulzmann, and M. Wehr. Type inference with constrained types. *tapos*, 5(1):35–55, 1999.
- [10] S. Peyton Jones. Proposal: TypeDirectedNameResolution. <https://ghc.haskell.org/trac/haskell-prime/wiki/TypeDirectedNameResolution>. Last modified: 2009-11-09, Accessed: 2015-03-25.
- [11] F. Pottier and D. Rémy. The essence of ML type inference. In B. C. Pierce, editor, *Advanced Topics in Types and Programming Languages*, chapter 10, pages 389–489. MIT Press, 2005.
- [12] D. Rémy. Typing record concatenation for free. In *POPL*, pages 166–176, 1992.
- [13] T. Schrijvers, S. L. P. Jones, M. Sulzmann, and D. Vytiniotis. Complete and decidable type inference for gadts. In *ICFP*, pages 341–352, 2009.
- [14] V. Simonet and F. Pottier. A constraint-based approach to guarded algebraic data types. *TOPLAS*, 29(1), 2007.
- [15] V. Trifonov and S. F. Smith. Subtyping constrained types. In *SAS*, pages 349–365, 1996.
- [16] D. Vytiniotis, S. L. P. Jones, T. Schrijvers, and M. Sulzmann. Outsidein(x) modular type inference with local assumptions. *J. Funct. Program.*, 21(4-5):333–412, 2011.
- [17] M. Wand. A simple algorithm and proof for type inference. *Fundamenta Informaticae*, 10:115–122, 1987.

A. Proofs

A.1 Proof of property † in Section 6.3

In this proof, given a triple $(\mathcal{G}, F^{<:}, \psi)$, we denote $(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}}$ the triple $(\mathcal{G}, \leq_{F^{<:}}, \psi)$.

We want to prove the following property about the constraint resolution system presented in Figure 14:

$$\begin{aligned} & \forall (C_1, \mathcal{G}_1, F^{<:}_1, \psi_1), (C_2, \mathcal{G}_2, F^{<:}_2, \psi_2), \\ & (C_1, \mathcal{G}_1, F^{<:}_1, \psi_1) \longrightarrow (C_2, \mathcal{G}_2, F^{<:}_2, \psi_2) \Rightarrow \\ & \forall \sigma, \sigma(\mathcal{G}_2, F^{<:}_2, \psi_2)^{\mathcal{M}} \models \sigma(C_2) \Rightarrow \\ & \exists \sigma', \sigma'(\mathcal{G}_1, F^{<:}_1, \psi_1)^{\mathcal{M}} \models \sigma'(C_1) \quad (\dagger) \end{aligned}$$

Proof. We prove this property by case analysis on the reduction:

$$(C_1, \mathcal{G}_1, F^{<:}_1, \psi_1) \longrightarrow (C_2, \mathcal{G}_2, F^{<:}_2, \psi_2)$$

- S-RESOLVE Assume:

$$(x^R \mapsto \delta \wedge C, \mathcal{G}, F^{<:}, \psi) \longrightarrow [\delta \mapsto x^D](C, \mathcal{G}, F^{<:}, \psi)$$

where $x^D \in R_{\mathcal{G}}(x^R)$ and let σ' be $[\delta \mapsto x^D]$.

Assume there is σ such that

$$\sigma(\sigma'(\mathcal{G}, F^{<:}, \psi))^{\mathcal{M}} \models \sigma(\sigma' C) \quad (\text{H})$$

then we want to prove:

$$\exists \sigma_1, \sigma_1(\mathcal{G}, F^{<:}, \sigma' \psi)^{\mathcal{M}} \models \sigma_1(x^R \mapsto \delta \wedge C)$$

We have:

1. $\vdash_{\mathcal{G}} x^R \mapsto x^D$ by correctness of the name resolution algorithm $R_{\mathcal{G}}()$,
2. $\vdash_{\sigma \sigma'} x^R \mapsto x^D$ by definition,
3. $\sigma \sigma'(\mathcal{G}, F^{<:}, \sigma' \psi)^{\mathcal{M}} \models \sigma \sigma' x^R \mapsto \delta$
4. $(\sigma \sigma')(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models (\sigma \sigma') C$ using H
5. $(\sigma \sigma')(\mathcal{G}, F^{<:}, \sigma' \psi)^{\mathcal{M}} \models (\sigma \sigma') C$ since $\sigma' \sigma' = \sigma'$
6. we conclude with $\sigma_1 = (\sigma \sigma')$ by C-AND rule of the constraint interpretation with 3. and 5.

- S-ASSOC Assume:

$$(x^D \rightsquigarrow_{\zeta} \wedge C, \mathcal{G}, F^{<:}, \psi) \longrightarrow [\zeta \mapsto S](C, \mathcal{G}, F^{<:}, \psi)$$

where $\mathcal{DSC}(x^D) = S$ and let σ' be $[\zeta \mapsto S]$.

Assume there is σ such that

$$\sigma(\sigma'(\mathcal{G}, F^{<:}, \psi))^{\mathcal{M}} \models \sigma(\sigma' C) \quad (\text{H})$$

then we want to prove:

$$\exists \sigma_1, \sigma_1(\mathcal{G}, F^{<:}, \sigma' \psi)^{\mathcal{M}} \models \sigma_1(x^D \rightsquigarrow_{\zeta} \wedge C)$$

We have:

1. $\mathcal{DSC}(x^D) = S$ by the rewriting rule condition
2. $\sigma \sigma'(\mathcal{G}, F^{<:}, \sigma' \psi)^{\mathcal{M}} \models \sigma \sigma' x^D \rightsquigarrow_{\zeta}$
3. $(\sigma \sigma')(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models (\sigma \sigma') C$ using H
4. $(\sigma \sigma')(\mathcal{G}, F^{<:}, \sigma' \psi)^{\mathcal{M}} \models (\sigma \sigma') C$ since $\sigma' \sigma' = \sigma'$
5. we conclude with $\sigma_1 = (\sigma \sigma')$ by C-AND rule of the constraint interpretation with 2. and 4.

- S-EQUAL Assume:

$$(T_1 \equiv T_2 \wedge C, \mathcal{G}, F^{<:}, \psi) \longrightarrow \sigma'(C, \mathcal{G}, F^{<:}, \psi)$$

where $\sigma' = \mathcal{U}(T_1, T_2)$.

Assume there is σ such that

$$\sigma(\sigma'(\mathcal{G}, F^{<:}, \psi))^{\mathcal{M}} \models \sigma(\sigma' C) \quad (\text{H})$$

then we want to prove:

$$\exists \sigma_1, \sigma_1(\mathcal{G}, F^{<:}, \sigma' \psi)^{\mathcal{M}} \models \sigma_1(T_1 \equiv T_2 \wedge C)$$

We have:

1. $\sigma' t_1 = \sigma' t_2$ by unification property
2. $\sigma \sigma'(\mathcal{G}, F^{<:}, \sigma' \psi)^{\mathcal{M}} \models \sigma \sigma' T_1 \equiv T_2$ by C-EQUAL rule and 1.
3. $(\sigma \sigma')(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models (\sigma \sigma') C$ using H
4. $(\sigma \sigma')(\mathcal{G}, F^{<:}, \sigma' \psi)^{\mathcal{M}} \models (\sigma \sigma') C$ since $\sigma' \sigma' = \sigma'$ by unification property
5. we conclude by C-AND rule of the constraint interpretation with 2. and 4.

- S-LUB Assume:

$$(T \text{ is } t_1 \sqcup t_2 \wedge C, \mathcal{G}, F^{<:}, \psi) \longrightarrow (T \equiv t \wedge C, \mathcal{G}, F^{<:}, \psi)$$

where $(t_1 \sqcup t_2) \xrightarrow{F^{<:}} t$.

Assume there is σ such that

$$\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma(t \equiv T \wedge C) \quad (\text{H})$$

then we want to prove:

$$\exists \sigma_1, \sigma_1(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma_1(T \text{ is } t_1 \sqcup t_2 \wedge C)$$

We have:

1. $\sigma t = \sigma T$ by inversion of C-AND and C-EQUAL semantics rules on H
2. $t = \sqcup_{F^{<:}} \{t_1, t_2\}$ by correctness of (\sqcup) reduction
3. $\sigma t = t \wedge \sigma t_1 = t_1 \wedge \sigma t_2 = t_2$ since these are ground terms
4. $\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma(T \text{ is } t_1 \sqcup t_2)$ since $F^{<:}$ is ground
5. $\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma(C)$ using H
6. we conclude by C-AND rule of the constraint interpretation with 4. and 5.

$(Int \sqcup Int) \xrightarrow{F^{<:}} Int$	(LUB-INT)
$(Rec(d) \sqcup Rec(d')) \xrightarrow{F^{<:}} \begin{cases} Rec(d) & \text{if } Rec(d') \leq_{F^{<:}} Rec(d) \\ (Rec(d) \sqcup Rec(d')) & \text{if } Rec(d') <: Rec(d') \in F^{<:} \end{cases}$	(LUB-REC)
$(Fun[t_1, t_2] \sqcup Fun[t_3, t_4]) \xrightarrow{F^{<:}} Fun[(t_1 \sqcup t_3), (t_2 \sqcup t_4)]$	(LUB-ARROW)
$(Int \sqcap Int) \xrightarrow{F^{<:}} Int$	(GLB-INT)
$(Rec(d) \sqcap Rec(d')) \xrightarrow{F^{<:}} \begin{cases} Rec(d) & \text{if } Rec(d) \leq_{F^{<:}} Rec(d') \\ Rec(d') & \text{if } Rec(d') \leq_{F^{<:}} Rec(d) \end{cases}$	(GLB-REC)
$(Fun[t_1, t_2] \sqcap Fun[t_3, t_4]) \xrightarrow{F^{<:}} Fun[(t_1 \sqcap t_3), (t_2 \sqcap t_4)]$	(GLB-ARROW)

Figure 15. LMR specific functions

- S-SUBTYPE Assume:

$$(t_1 \leq_{F^{<:}} t_2 \wedge C, \mathcal{G}, F^{<:}, \psi) \longrightarrow (C, \mathcal{G}, F^{<:}, \psi)$$

Assume there is σ such that

$$\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma(C) \quad (\text{H})$$

then we want to prove:

$$\exists \sigma_1, \sigma_1(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma_1(t_1 \leq_{F^{<:}} t_2 \wedge C)$$

We have:

1. $t_1 \leq_{F^{<:}} t_2$ by reduction rule hypothesis
2. $\sigma t_1 \leq_{\sigma(F^{<:})} \sigma t_2$ by \leq_x monotonicity and since t_1 and t_2 are ground
3. $\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma(t_1 \leq_{F^{<:}} t_2)$ by C-Subtype semantics rule
4. we conclude by C-AND rule of the constraint interpretation with 3. and H

- S-DECLTYPEFIRST Assume:

$$(x^D : T \wedge C, \mathcal{G}, F^{<:}, \psi) \longrightarrow (C, \mathcal{G}, F^{<:}, \{x^D \mapsto T\} \cup \psi)$$

Assume there is σ such that

$$\sigma(\mathcal{G}, F^{<:}, \{x^D \mapsto T\} \cup \psi)^{\mathcal{M}} \models \sigma(C) \quad (\text{H})$$

then we want to prove:

$$\exists \sigma_1, \sigma_1(\mathcal{G}, F^{<:}, \{x^D \mapsto T\} \cup \psi)^{\mathcal{M}} \models \sigma_1(x^D : T \wedge C)$$

We have:

1. $\sigma(\mathcal{G}, F^{<:}, \{x^D \mapsto T\} \cup \psi)^{\mathcal{M}} \models \sigma(x^D : T)$ by C-TypeOf semantics rule
2. we conclude by C-AND rule of the constraint interpretation with 1. and H

- S-DECLTYPENEXT Assume:

$$(x^D : T \wedge C, \mathcal{G}, F^{<:}, \psi) \longrightarrow (\psi(x^D) \equiv T \wedge C, \mathcal{G}, F^{<:}, \psi)$$

Assume there is σ such that

$$\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma(\psi(x^D) \equiv T \wedge C) \quad (\text{H})$$

then we want to prove:

$$\exists \sigma_1, \sigma_1(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma_1(x^D : T \wedge C)$$

We have:

1. $\sigma \psi(x^D) = \sigma T$ by inversion of C-AND and C-Equal semantics rules
2. $\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma(x^D : T)$ by C-TypeOf rule
3. $\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma(C)$ using H
4. we conclude by C-AND rule of the constraint interpretation with 2. and 3.

- S-TRUE Assume:

$$(True \wedge C, \mathcal{G}, F^{<:}, \psi) \longrightarrow (C, \mathcal{G}, F^{<:}, \psi)$$

Assume there is σ such that:

$$\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma(C) \quad (\text{H})$$

then we have:

1. $\sigma(\mathcal{G}, F^{<:}, \psi)^{\mathcal{M}} \models \sigma True$ by C-True rule
2. we conclude by C-AND rule of the constraint interpretation with 1. and H.

□

B. LMR Least upper bound computation

Algorithm in Fig. 15 is the least upper bound computation for the LMR language. The correctness on this algorithm relies on the property that each *Rec* type has a unique direct ancestor.

TUD-SERG-2015-006
ISSN 1872-5392

