Type Inference:
Unification, Parameterized Types,
Underconstrained Systems, and Principal Types

Shriram Krishnamurthi

2003-11-10

1 Solving Type Constraints

To solve type constraints, we turn to a classic algorithm: *unification*. Unification consumes a set of constraints and either

- identifies inconsistencies amongst the constraints, or

- generates a *substitution* that represents the solution of the constraints.

A substitution is a mapping from names to constants. In our universe, inconsistencies indicate type errors, the constants are terms in the type language (such as number and number $\rightarrow$ boolean), and the names are identifiers representing the values of individual boxes (thus $\boxed{4}$ is a funny notation for an identifier that represents the type of the expression labeled 4).

The unification algorithm is extremely simple. Begin with an empty substitution. Push all the constraints onto a stack. If the stack is empty, return the substitution; otherwise, pop the constraint $X = Y$ off the stack:

1. If X and Y are identical constants, do nothing.

2. If X and Y are identical identifiers, do nothing.

3. If X is an identifier, replace all occurrences of X by Y both on the stack and in the substitution, and add $X \mapsto Y$ to the substitution.

4. If Y is an identifier, replace all occurrences of Y by X both on the stack and in the substitution, and add $Y \mapsto X$ to the substitution.

5. If X is of the form $C(X_1, \ldots, X_n)$ for some constructor C,[1] and Y is of the form $C(Y_1, \ldots, Y_n)$ (i.e., it has the same constructor), then push $X_i = Y_i$ for all $1 \leq i \leq n$ onto the stack.

6. Otherwise, X and Y do not unify. Report an error.

Does this this algorithm terminate? On every iteration of the main loop, it pops a constraint off the stack. In some cases, however, we push new constraints on. The *size* of each of these constraints is, however, smaller than the constraint just popped. Therefore, the total number of iterations cannot be greater than the sum of the sizes of the initial constraint set. The stack must therefore eventually become empty.

Problem What is this algorithm's complexity?

---

[1]Our type language has one built-in type constructor, namely $\rightarrow$. Parameterized dataypes, such as lists, introduce additional type constructors.

2 Example of Unification at Work

Let's consider the following example:

$\boxed{1}(\boxed{2}\textbf{(lambda } (x)\ x)$
$\quad \boxed{3}7)$

This generates the following constraints:

$$\llbracket \boxed{2} \rrbracket = \llbracket \boxed{3} \rrbracket \rightarrow \llbracket \boxed{1} \rrbracket$$

$$\llbracket \boxed{2} \rrbracket = \llbracket x \rrbracket \rightarrow \llbracket x \rrbracket$$

$$\llbracket \boxed{3} \rrbracket = \mathsf{number}$$

The unification algorithm works as follows:

Action	Stack	Substitution
Initialize	$\llbracket \boxed{2} \rrbracket = \llbracket \boxed{3} \rrbracket \rightarrow \llbracket \boxed{1} \rrbracket$ $\llbracket \boxed{2} \rrbracket = \llbracket x \rrbracket \rightarrow \llbracket x \rrbracket$ $\llbracket \boxed{3} \rrbracket = \mathsf{number}$	empty
Step 3	$\llbracket \boxed{3} \rrbracket \rightarrow \llbracket \boxed{1} \rrbracket = \llbracket x \rrbracket \rightarrow \llbracket x \rrbracket$ $\llbracket \boxed{3} \rrbracket = \mathsf{number}$	$\llbracket \boxed{2} \rrbracket \mapsto \llbracket \boxed{3} \rrbracket \rightarrow \llbracket \boxed{1} \rrbracket$
Step 5	$\llbracket \boxed{3} \rrbracket = \llbracket x \rrbracket$ $\llbracket \boxed{1} \rrbracket = \llbracket x \rrbracket$ $\llbracket \boxed{3} \rrbracket = \mathsf{number}$	$\llbracket \boxed{2} \rrbracket \mapsto \llbracket \boxed{3} \rrbracket \rightarrow \llbracket \boxed{1} \rrbracket$
Step 3	$\llbracket \boxed{1} \rrbracket = \llbracket x \rrbracket$ $\llbracket x \rrbracket = \mathsf{number}$	$\llbracket \boxed{2} \rrbracket \mapsto \llbracket x \rrbracket \rightarrow \llbracket \boxed{1} \rrbracket$ $\llbracket \boxed{3} \rrbracket \mapsto \llbracket x \rrbracket$
Step 3	$\llbracket x \rrbracket = \mathsf{number}$	$\llbracket \boxed{2} \rrbracket \mapsto \llbracket x \rrbracket \rightarrow \llbracket x \rrbracket$ $\llbracket \boxed{3} \rrbracket \mapsto \llbracket x \rrbracket$ $\llbracket \boxed{1} \rrbracket \mapsto \llbracket x \rrbracket$
Step 3	empty	$\llbracket \boxed{2} \rrbracket \mapsto \mathsf{number} \rightarrow \mathsf{number}$ $\llbracket \boxed{3} \rrbracket \mapsto \mathsf{number}$ $\llbracket \boxed{1} \rrbracket \mapsto \mathsf{number}$ $\llbracket x \rrbracket \mapsto \mathsf{number}$

At this point, we have solutions for all the sub-expressions *and* we know that the constraint set is consistent.

Writing these in detail in notes format is painstaking, but it's usually easy to do them on paper by just crossing out old values when performing a substitution. Be sure to work through our examples for more practice with unification!

3 Parameterized Types

In the presentation of unification above, we saw only one kind of type constructor, for functions: $\rightarrow$. A regular programming language will typically have many more constuctors. A common source of parameterized types is *containers* such as lists, trees, queues, and so forth. Thus, for instance, it is common to think of lists as parameterized over their content, thus yielding list(number), list(symbol), list(list(number)), list(number $\rightarrow$ symbol) and so on. Identifying list as one of the type constructors for the unification algorithm suffices for typing *unityped* lists.

4 The "Occurs" Check

Suppose we generate the following type constraint:

$$\mathsf{list}(\llbracket x \rrbracket) = \mathsf{list}(\mathsf{list}(\llbracket x \rrbracket))$$

By Step 5, we should push $[\![x]\!] = \mathsf{list}([\![x]\!])$ onto the stack. Eventually, obeying to Step 3, we will add the mapping $[\![x]\!] \mapsto \mathsf{list}([\![x]\!])$ to the substition but, in the process, attempt to replace all instances of $[\![x]\!]$ in the substitution with the right-hand side, which does not terminate.

This is a familiar problem: we saw it earlier when trying to define substitution in the presence of recursion. Because these are problematic to handle, a traditional unification algorithm checks in Steps 3 and 4 whether the identifier about to be (re)bound in the substitution *occurs* in the term that will take its place. If the identifier does occur, the unifier halts with an error.[2] Otherwise, the algorithm proceeds as before.

Puzzle   Write a program that will generate the above constraint!

5   Underconstrained Systems

We have seen in our previous lecture that if the system has too many competing constraints—for instance, forcing an identifier to have both type number and $\mathsf{boolean}$—there can be no satisfying type assignment, so the system should halt with an error. We saw the indication of this error informally in the previous lecture. Working through that example with the unification algorithm should confirm that the algorithm does indeed detect the error (in Step 6).

But what if the system is *under*-constrained? This is interesting, because some of the program identifiers never get assigned a type! In a procedure such as *map*, for instance:

(**define** (*map f l*)
  (**cond**
    [(*empty? l*) *empty*]
    [(*cons? l*) (*cons* (*f* (*first l*))
                (*map f* (*rest l*)))]))

Working through the constraint generation and unification algorithm for examples such as this is educational. Doing so reveals that there is no constraint on the type passed as a parameter to the list type constructor. Working through the steps, we get a type for *map* of this form:

$$(\alpha{\to}\beta) \times \mathsf{list}(\alpha){\to}\mathsf{list}(\beta)$$

where α and β are unconstrained type identifiers. This is the same type we obtained through explicit parametric polymorphism. . . except that the unification algorithm has found it for us *automatically*!

6   Principal Types

The type generated by this Hindley-Milner[3] system has a particularly pleasing property: it is a *principal* type. A principal type has the following property. For a term t, consider a type τ. τ is a principal type of t if, for any other type τ' that types t, there exists a substitution (perhaps empty) that, when applied to τ, yields τ'.

There are a few ways of re-phrasing the above:

- The Hindley-Milner type system infers the "most general" type for a term.

- The type generated by the Hindler-Milner type system imposes fewest constraints on the program's behavior. In particular, it imposes constraints necessary for type soundness, but no more.

From a software engineering perspective, this is very attractive: it means a programmer could not possibly annotate a procedure with a more general type than the type inference algorithm would derive. Thus, using the algorithm instead of performing manual annotation will not restrict the reusability of code, and may even increase it (because the programmer's annotation may mistakenly overconstrain the type). Of course, there are other good reasons for manual annotation, such as documentation and readability, so a good programming style will mix annotation and inference judiciously.

---

[2]This is not the only reasonable behavior! It is possible to define *fixed-point types*, which are solutions to the circular constraint equations above. This topic is, however, beyond the scope of this course.

[3]Named for Roger Hindley and Robin Milner, who independently discovered such a type system in the late 1960s and early 1970s.