

Scalable Component Abstractions

Martin Odersky

Swiss Federal Institute of Technology Lausanne (EPFL)

(joint work with Matthias Zenger, Google)

Component Software – State of the Art

In *principle*, software should be constructed from re-usable parts (“components”).

In *practice*, software is still most often written “from scratch”, more like a craft than an industry.

Programming languages share part of the blame for this.

Most existing languages offer only limited support for components.

This holds in particular for statically typed languages such as Java and C#.

How To Do Better?

Hypothesis 1: Languages for components need to be *scalable*; the same concepts should describe small as well as large parts.

Hypothesis 2: Scalability can be provided by unifying and generalizing functional and object-oriented programming concepts.

To validate these hypotheses we have designed and implemented a concrete programming language, **Scala**.

Part I: A Quick Introduction to Scala

Some key aspects of Scala are:

- (1) interoperability with Java and .NET,
- (2) a uniform object model,
- (3) higher-order functions,
- (4) uniform abstraction concepts for both types and values,
- (5) symmetric mixins for composing classes.
- (6) object decomposition with pattern matching,
- (7) XML support.

(1) – (5) are quickly explained in the following.

1. A Java Like Language

Here is a sample program in Java:

```
class PrintOptions {
    public static void main(String[] args) {
        System.out.println(" Options selected:");
        for (int i = 0; i < args.length; i++)
            if (args[i].startsWith(" -"))
                System.out.println(" " +args[i].substring(1));
    }
}
```

And here is the same program in Scala:

```
object PrintOptions {
    def main(args: Array[String]): unit = {
        System.out.println(" Options selected:");
        for (val arg ← args)
            if (arg.startsWith(" -"))
                System.out.println(" " +arg.substring(1));
    }
}
```

Interoperability

Scala is completely interoperable with Java (and more recently also to C#).

A Scala component can:

- access all methods and fields of a Java component,
- create instances of Java classes,
- inherit from Java classes and implement Java interfaces,
- be itself instantiated and called from a Java component.

None of this requires glue code or special tools.

This makes it very easy to mix Scala and Java components in one application.

2. A Unified Object Model

In Scala, every value is an object and every operation is a method invocation.

Example: A class for natural numbers

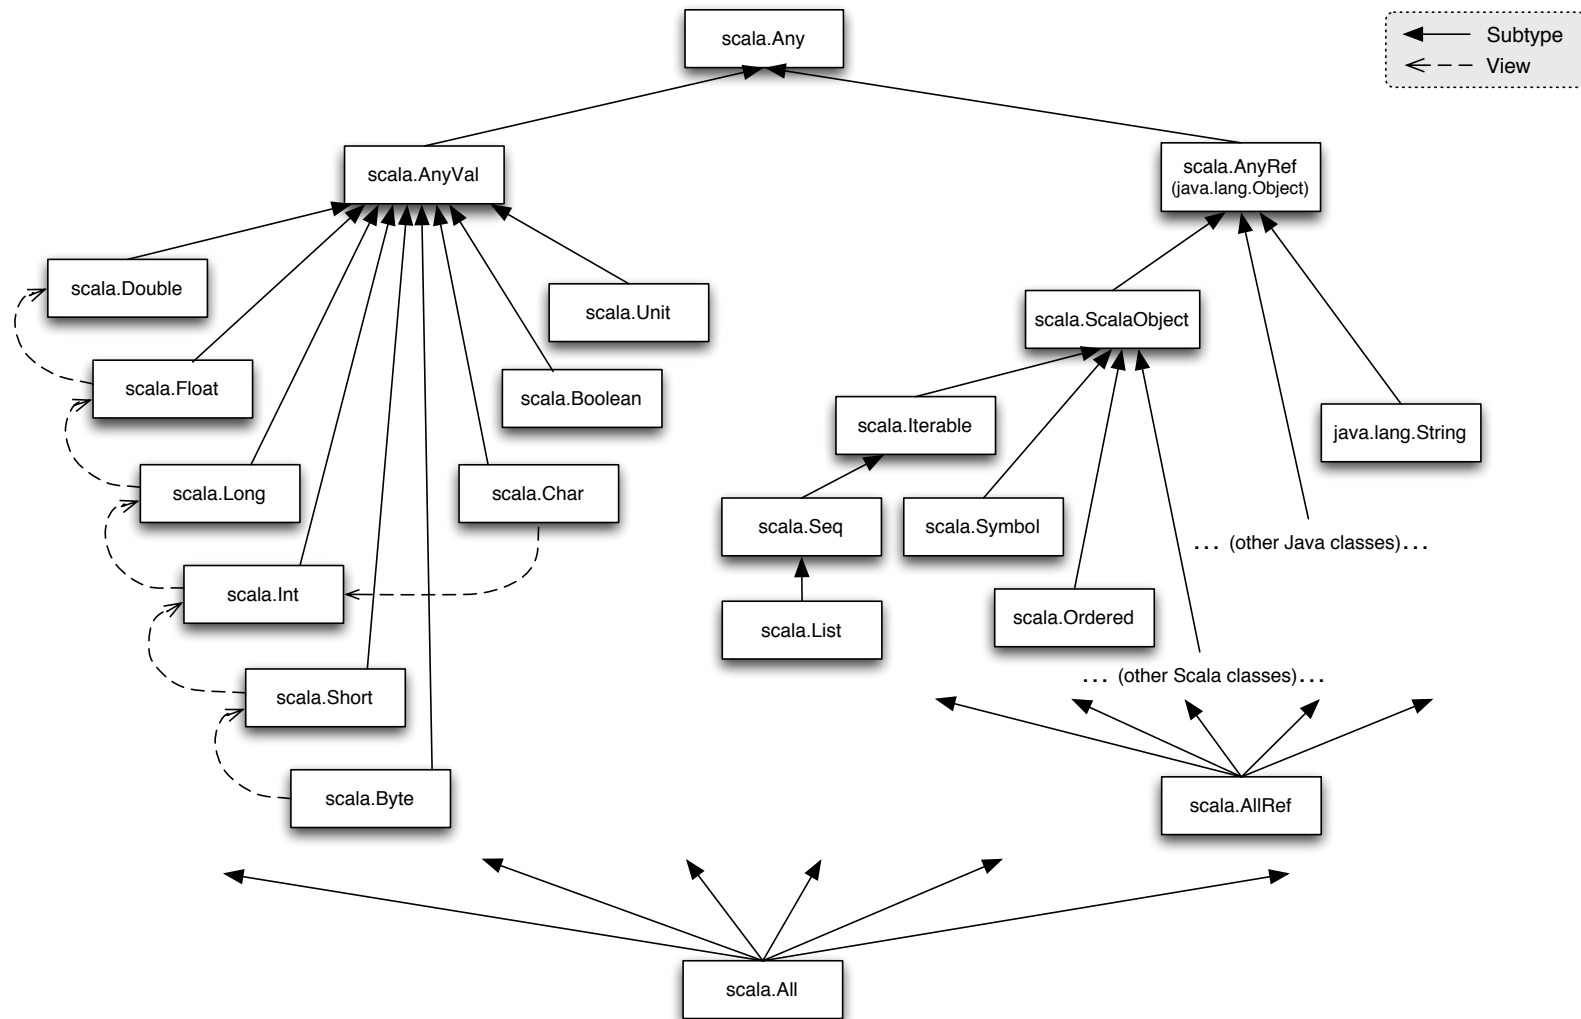
```
abstract class Nat {  
  def isZero: boolean;  
  def pred: Nat;  
  def succ: Nat = new Succ(this);  
  def + (x: Nat): Nat = if (x.isZero) this else succ + x.pred;  
  def - (x: Nat): Nat = if (x.isZero) this else pred - x.pred;  
}
```

Here are the two canonical implementations of Nat:

```
class Succ(n: Nat) extends Nat {  
  def isZero: boolean = false;  
  def pred: Nat = n  
}
```

```
object Zero extends Nat {  
  def isZero: boolean = true;  
  def pred: Nat =  
    throw new Error("Zero.pred");  
}
```

Scala's Class Hierarchy



3. Operations Are Objects

- Scala is a functional language, in the sense that every function is a value.
- Functions can be anonymous, curried, or nested inside each other.
- Familiar higher-order functions are implemented as methods of Scala classes. E.g.:
matrix exists (row \Rightarrow row forall (0 ==)))
- Here, matrix could be of type of List[List[int]], using Scala's List class:

```
class List[+T] {  
  def isEmpty: boolean;  
  def head: T  
  def tail: List[T];  
  def exists(p: T  $\Rightarrow$  boolean): boolean =  
    !isEmpty && (p(head) || (tail exists p));  
  ... }  
}
```

Functions are Objects

- If functions are values, and values are objects, it follows that functions themselves are objects.

- In fact, the function type $S \Rightarrow T$ is equivalent to `scala.Function1[S, T]`

where `Function1` is defined as follows in the standard Scala library:

```
abstract class Function1[-S, +T] { def apply(x: S): T }
```

(Analogous conventions exist for functions with more than one argument.)

- Hence, functions are interpreted as objects with `apply` methods.
- For example, the anonymous “incrementer” function `x: int \Rightarrow x + 1` is expanded as follows.

```
new Function1[int, int] { def apply(x: int): int = x + 1 }
```

Part II: Components

A *component* is a program part, to be combined with other parts in larger applications.

Requirement: Components should be *reusable*.

To be reusable in new contexts, a component needs *interfaces* describing its *provided* as well as its *required* services.

Most current components are not very reusable.

Most current languages can specify only provided services, not required services.

Note: **Component \neq API !**

No Statics!

A component should refer to other components not by hard links, but only through its required interfaces.

Another way of expressing this is:

All references of a component to others should be via its members or parameters.

In particular, there should be no global static data or methods that are directly accessed by other components.

This principle is not new (cf: the *Law of Demeter*)

But it is surprisingly difficult to achieve.

Functors

One established language abstraction for components are SML functors.

Here,

Component $\hat{=}$ *Functor or Structure*

Interface $\hat{=}$ *Signature*

Required Component $\hat{=}$ *Functor Parameter*

Composition $\hat{=}$ *Functor Application*

Sub-components are identified via sharing constraints.

Shortcomings:

- No recursive references between components
- No inheritance with overriding
- Structures are not first class.

Components in Scala

In Scala:

<i>Component</i>	$\hat{=}$	<i>Class</i>
<i>Interface</i>	$\hat{=}$	<i>Abstract Class</i>
<i>Required Component</i>	$\hat{=}$	<i>Abstract Member</i> or “Self”
<i>Composition</i>	$\hat{=}$	<i>Symmetric Mixin Composition</i>

Advantages:

- Components instantiate to objects, which are first-class values.
- Recursive references between components are supported.
- Inheritance with overriding is supported.
- Sub-components are identified by name
⇒ no explicit “wiring” is needed.

Language Constructs for Components

Scala has three concepts which are particularly interesting in component systems.

- *Abstract type members* allow to abstract over types that are members of objects.
- *Selftype annotations* allow to abstract over the type of “self”.
- *Symmetric mixin composition* provides a flexible way to compose components and component types.

All three abstractions have their theoretical foundation in the νObj calculus [Odersky et al., ECOOP03].

They subsume SML modules.

More precisely, (generative) SML modules can be encoded in νObj , but not *vice versa*.

Component Abstraction

There are two principal forms of abstraction in programming languages:

parameterization (functional)

abstract members (object-oriented)

Scala supports both styles of abstraction for types as well as values.

Both types and values can be parameters, and both can be abstract members.

Abstract Types

Here is a type of “cells” using object-oriented abstraction.

```
abstract class AbsCell {  
  type T;  
  val init: T;  
  private var value: T = init;  
  def get: T = value;  
  def set(x: T): unit = { value = x }  
}
```

The AbsCell class has an abstract type member T and an abstract value member init.

Instances of that class can be created by implementing these abstract members with concrete definitions.

```
val cell = new AbsCell { type T = int; val init = 1 }  
cell.set(cell.get * 2)
```

The type of cell is AbsCell { type T = int }.

Path-dependent Types

It is also possible to access `AbsCell` without knowing the binding of its type member.

For instance: `def reset(c: AbsCell): unit = c.set(c.init);`

Why does this work?

- `c.init` has type `c.T`
- The method `c.set` has type `c.T ⇒ unit`.
- So the formal parameter type and the argument type coincide.

`c.T` is an instance of a *path-dependent* type.

In general, such a type has the form $x_0. \dots .x_n.t$, where $n \geq 0$,

- x_0 is an immutable value
- x_1, \dots, x_n are immutable fields, and
- t is a type member of x_n .

Safety Requirement

Path-dependent types rely on the immutability of the prefix path.

Here is an example where immutability is violated.

```
var flip = false;
def f(): AbsCell = {
  flip = !flip;
  if (flip) new AbsCell { type T = int; val init = 1 }
  else new AbsCell { type T = String; val init = "" }
}
f().set(f().get) // illegal!
```

Scala's type system does not admit the last statement, because the computed type of `f().get` would be `f().T`.

This type is not well-formed, since the method call `f()` is not a path.

Family Polymorphism

Scala's abstract type concept is particularly well suited for *family polymorphism*, where several types vary together covariantly.

Example: The subject/observer pattern (also known as *publish/subscribe*):

```
abstract class SubjectObserver {
  type S <: Subject;
  type O <: Observer;
  abstract class Subject : S {
    private var observers: List[O] = List();
    def subscribe(obs: O) = observers = obs :: observers;
    def publish = for (val obs ← observers) obs.notify(this);
  }
  abstract class Observer { def notify(sub: S): unit; }
}
```

The top-level class SubjectObserver has two member classes:

- The Subject class defines methods subscribe and publish.
- The Observer class only declares an abstract method notify.

Note that the Subject and Observer classes do not directly refer to each other.

Instead, they refer to two abstract types S and O which are bounded by Subject and Observer.

Self-Types

Note that class `Subject` carries an explicit type annotation:

```
class Subject : S { ...
```

`S` is called a *self-type* of class `Subject`. If a self-type is given, it is taken as the type of `this` inside the class.

In class `Subject`, the self-type is necessary to render the call `obs.notify(this)` type-correct.

Without an explicit type annotation, the self type is taken to be the type of the class itself.

Safety Requirement

Self types can be arbitrary; they need not have a relation with the class being defined.

Type soundness is still guaranteed, because of two requirements:

1. The self-type of a class must be a subtype of the self-types of all its base classes.
2. When instantiating a class in a **new** expression, it is checked that the self type of the class is a supertype of the type of the object being created.

[In a sense, explicit self types abstract over **this** (and are thus similar to abstract members). They represent a constraint on the type of **this** which is to be fulfilled by subclasses.]

Family Polymorphism ctd

The mechanism defined in the publish/subscribe pattern can be used by inheriting from SubjectObserver

Example:

```
object SensorReader extends SubjectObserver {
  type S = Sensor;
  type O = Display;
  abstract class Sensor extends Subject {
    val label: String;
    var value: double = 0.0;
    def changeValue(v: double) = { value = v; publish; }
  }
  abstract class Display extends Observer {
    def println(s: String) = ...
    def notify(sub: Sensor) = println(sub.label + " has value " + sub.value);
  }
}
```


In object `SensorReader`, type `S` is bound to `Sensor`, and type `O` is bound to `Display`.

The two formerly abstract types are now defined by overriding definitions.

This “tying the knot” is always necessary when creating a concrete class instance.

On the other hand, it would also have been possible to define an abstract `SensorReader` class which could be refined further by client code.

In this case, the two abstract types would have been overridden again by abstract type definitions.

```
abstract class AbsSensorReader extends SubjectObserver {  
    type S <: Sensor;  
    type O <: Display;  
    ...  
}
```

Example: Symbol Tables

Here's an example, which reflects a learning curve I had when writing extensible compiler components.

- Compilers need to model symbols and types.
- Each aspect depends on the other.
- Both aspects require substantial pieces of code.

The first attempt of writing a Scala compiler in Scala defined two global objects, one for each aspect:

First Attempt: Global Data

```
object Symbols {  
  class Symbol {  
    def tpe: Types.Type;  
    ...  
  }  
  // static data for symbols  
}
```

```
object Types {  
  class Type {  
    def sym: Symbols.Symbol  
    ...  
  }  
  // static data for types  
}
```

Problems:

1. Symbols and Types contain hard references to each other.
Hence, impossible to adapt one while keeping the other.
2. Symbols and Types contain static data.

Hence the compiler is not *reentrant*, multiple copies of it cannot run in the same OS process.

(This is a problem for a Scala Eclipse plugin, for instance).

Second Attempt: Nesting

Static data can be avoided by nesting the Symbols and Types objects in a common enclosing class:

```
class SymbolTable {  
  object Symbols {  
    class Symbol { def tpe: Types.Type; ... }  
  }  
  object Types {  
    class Type { def sym: Symbols.Symbol; ... }  
  }  
}
```

This solves the re-entrancy problem.

But it does not solve the component reuse problem.

- Symbols and Types still contain hard references to each other.
- Worse, since they are nested in an enclosing object they can no longer be written and compiled separately.

Third Attempt: A Component-Based Solution

Question: How can one express the required services of a component?

Answer: By abstracting over them!

Two forms of abstraction: *parameterization* and *abstract members*.

Only abstract members can express recursive dependencies, so we will use them.

```
abstract class Symbols {  
  type Type;  
  class Symbol { def tpe: Type }  
}
```

```
abstract class Types {  
  type Symbol;  
  class Type { def sym: Symbol }  
}
```

Symbols and Types are now classes that each abstract over the identity of the “other type”. How can they be combined?

Symmetric Mixin Composition

Here's how:

```
class SymbolTable extends Symbols with Types;
```

Instances of the `SymbolTable` class contain all members of `Symbols` as well as all members of `Types`.

Concrete definitions in either base class override abstract definitions in the other.

[Symmetric mixin composition generalizes the single inheritance + interfaces concept of Java and C#.

It is similar to *traits* [Schaerli et al, ECOOP 2003], but is more flexible since base classes may contain state.

]

Fourth Attempt: Mixins + Self Types

The last solution modeled required types by abstract types.

This is limiting, because one cannot instantiate or inherit an abstract type.

A more general approach also makes use of *self-types*:

```
class Symbols
  : Symbols with Types {
    class Symbol { def tpe: Type }
  }

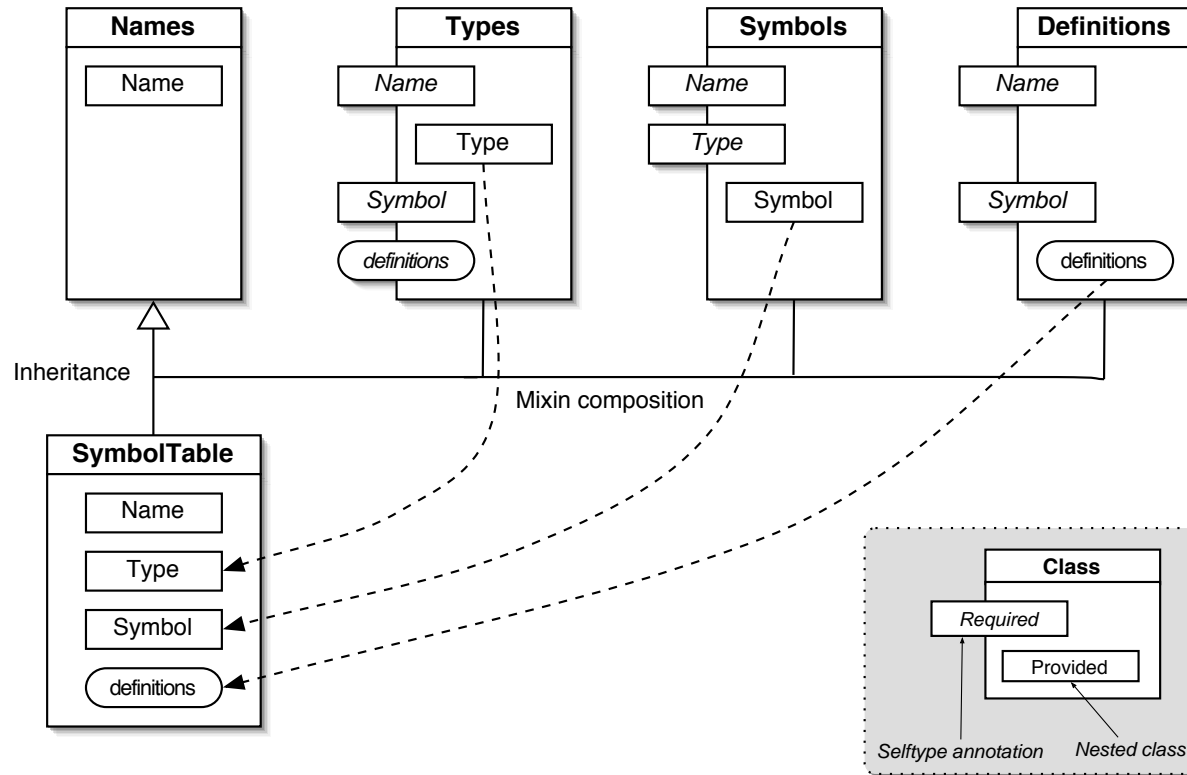
class Types
  : Types with Symbols {
    class Type { def sym: Symbol }
  }

class SymbolTable extends Symbols with Types;
```

Here, every component has a self type that contains all required components.

Symbol Table Schema

Here's a schematic drawing of *scalac*'s symbol table:



We see that besides **Symbols** and **Types** there are several other classes that also depend recursively on each other.

Benefits

1. The presented scheme is very *general* – any combination of static modules can be lifted to an assembly of components.
2. Components have *documented interfaces* for required as well as provided services.
3. Components can be *multiply instantiated*
 \Rightarrow *Reentrancy* is no problem.
4. Components can be flexibly *extended* and *adapted*.

Example: Logging

As an example of component adaptation, consider adding some logging facility to the compiler.

Say, we want a log of every symbol and type creation.

To print logging information, we use the following abstract class, which can be instantiated with arbitrary implementations.

```
abstract class Log {  
    def println(s: String): unit  
}
```

The problem is how insert calls to the `println` method into an existing compiler

- without changing source code,
- with clean separation of concerns,
- without using AOP.

Logging Classes

The idea is that the tester of the compiler would create subclasses of components which contain the logging code. E.g.

```
abstract class LogSymbols extends Symbols {  
  val log: Log;  
  override def newTermSymbol(name: Name): TermSymbol = {  
    val x = super.newTermSymbol(name);  
    log.println("creating term symbol " + name);  
    x  
  }  
  ...  
}
```

... and similarly for LogTypes.

How can these classes be integrated in the compiler?

Inserting Behavior by Mixin Composition

Here's an outline of the Scala compiler root class:

```
class ScalaCompiler extends SymbolTable with ... { ... }
```

To create a logging compiler, we extend this class as follows:

```
class TestCompiler extends ScalaCompiler with LogSymbols with LogTypes {  
  val log = new ConsoleLog;  
}
```

Now, every call to a factory method like `newTermSymbol` is re-interpreted as a call to the corresponding method in `LogSymbols`.

Note that the mixin-override is non-local – methods are overridden even if they are defined by indirectly inherited classes.

SelfTypes + Mixins *vs.* AOP

Similar strategies work for many adaptations for which aspect-oriented programming is usually proposed. E.g.

- security checks
- synchronization
- choices of data representation (e.g. sparse vs dense arrays)

Generally, one can handle all before/after advice on method join-points in this way.

Relationship between Scala and Other Languages

Main influences on the Scala design:

1. Java, C# for their syntax, basic types, and class libraries,
2. Smalltalk for its uniform object model,
3. Beta for systematic nesting,
4. ML, Haskell for many of the functional aspects.

(Too many influences in details to list them all)

Component Composition in Other Languages

1. Java, C#: Mostly static data. No true reusable components.
2. Smalltalk: Object-level instead of class-level composition. Less problems (and less security) for lack of static types.
3. Beta: Nesting + an extra-language layer of “fragments”.
4. ML: Parameterization only; No inheritance or recursive dependencies are possible.

Related Language Research

Mixin composition : Bracha (linear), Duggan, Hirschowitz
(mixin-modules), Schaerli et al. (traits), Flatt et al. (units,
Jiazzi), Zenger (Keris).

Abstract type members : Ernst (gbeta), Jolly et al. (Concord).

Explicit self types : Vuillon and Rémy (OCaml)

Conclusion

- Scala enables a new method for software construction “in the large”.
- Static data and references are replaced by member and selftype abstraction and symmetric mixin composition.
- Programs are classes which can be instantiated multiple times.
- This enables better separation of concerns and more flexible component adaptation.

Try it out: scala.epfl.ch

Thanks to the (past and present) members of the Scala team:

Philippe Altherr, Vincent Cremet, Julian Dragos, Burak Emir, Sebastian Maneth, Stéphane Micheloud, Nikolay Mihaylov, Michel Schinz, Erik Stenman, Matthias Zenger.

Addendum I: The Expression Problem

With similar abstraction techniques, we can solve the *expression problem*:

How can a system be extended at the same time with new data variants and with new operations over data?

Requirements:

1. Separate compilation,
2. strong static type safety,
3. no code modification.

See:

[1] Matthias Zenger and Martin Odersky. Independently Extensible Solutions to the Expression Problem. EPFL Technical Report IC/2004/33

Addendum II: Generalized Algebraic Data Types

Here's how GADT's would be expressed in Scala.

```
abstract class Term[T];
case class Lit(x: int);
case class Succ(t: Term[int]) extends Term[int];
case class IsZero(t: Term[int]) extends Term[boolean];
case class If[T](c: Term[boolean], t1: Term[T], t2: Term[T]) extends Term[T];

def eval[a](t: Term[a]): a = t match {
  case Lit(n)           => n
  case Succ(u)          => eval(u) + 1
  case IsZero(u)        => eval(u) == 0
  case If(c, u1, u2) => if (eval(c)) eval(u1) else eval(u2)
}
```

Caveat: In current Scala, this would not type check.

Reason: Type variable bindings in a pattern are “forgotten” on the

right of the “ \Rightarrow ”.

Hence, `eval`'s result type would be `Any`, not `a`.

But maybe we should change that?