

Smart Modules for Java-like Languages*

Davide Ancona, Giovanni Lagorio, and Elena Zucca

DISI - Università di Genova
Via Dodecaneso, 35, 16146 Genova (Italy)
email: {davide,lagorio,zucca}@disi.unige.it

Abstract. We present SMARTJAVAMOD, a language of mixin modules supporting *compositional compilation*, and constructed on top of the Java language. More in detail, this means that basic modules are collections of Java classes which can be typechecked in isolation, inferring constraints on missing classes and allowing safe reuse of the module in as many contexts as possible. Furthermore, it is possible to write structured module expressions by means of a set of module operators, and a type system at the module level ensures type safety, in the sense that we can always reduce a module expression to a well-formed collection of Java classes. What we obtain is a module language which is extremely flexible and allows the encoding (without any need of enriching the core level, that is, the Java language) of a variety of constructs supporting software reuse and extensibility.

1 Introduction

One of the main reasons of the success of the object-oriented paradigm is its good support for software reuse and extensibility; yet, mainstream object-oriented strongly typed languages still suffer from problems which prevent code to be easily and modularly reused in many circumstances.

Witness of this problem is the amount of recent papers which can be found in literature which aim at enhancing software extensibility in object-oriented languages [12, 3, 2, 15, 13, 8]. The proposed solutions are very disparate, ranging from the introduction of single new abstraction mechanisms that can be more or less easily integrated with existing languages, to more general and pervasive solutions leading to new languages and even paradigms [24, 23].

The approach followed here is partly different from both, and is based on the idea that software extensibility can be enhanced without extending a language in itself, but rather by building on top of this language a *module system*, as it has been successfully achieved in the functional paradigm (see SML for instance). Indeed, Java (and the other mainstream object-oriented languages) lack a real module system, partly because the notion of class is overloaded and, in fact, classes play the double role of object templates and modules. However, as firstly

* Partially supported by Dynamic Assembly, Reconfiguration and Type-checking - EC project IST-2001-33477, APPSEM II - Thematic network IST-2001-38957, and MIUR EOS - Extensible Object Systems.

recognized by Szyperski [30], classes do not offer all features expected from a modern module system (see also [17] for an extended discussion). Java packages are a symptom of this deficiency: they provide a better support than classes for name space management and code structuring, but still they are not real modules.

In particular, the module system we consider in this paper is based on the notion of *mixin module*, which has received a quite consolidated consensus as an effective means for enhancing software reuse. Mixin modules are a generalization, due to Cook and Bracha [10, 11], of the notion of *mixin class* (their class parametric in the parent) [9, 20, 2]. With respect to conventional modules, mixin modules offer advanced features including *mutually recursive* modules and *virtual* module components which can be redefined via an *overriding* operator. Mixin modules are equipped with a quite powerful set of operations which allow flexible manipulation and reuse of code in a type safe style, and they offer general and uniform solutions to the problem of software extension, so that some other mechanisms (as, for instance, traits, generic classes, mixin classes) can be, totally or partially, recovered in them.

While the theoretical foundations of mixin modules have been already investigated in details [4], much work still needs to be done for integrating real languages with systems of mixin modules. In particular, designing a system of mixin modules on top of Java is not a simple task. One of the main challenges is to devise a type system expressive enough for supporting separate typechecking and compilation of modules, without losing the potentialities offered by mixins for reusing and extending code. In [3], we described a proposal, called JAVAMOD (supporting separate typechecking via module interfaces, and module expressions constructed by *merge*, *renaming* and *hiding* operators), which, however, had many limitations. In particular, the programmer had to explicitly annotate with their required types the classes used as parameters in modules, thus making module declarations cumbersome; the type system was too restrictive (see below); furthermore, virtual classes were not supported.

In this paper we present SMARTJAVAMOD, a major revision of the design given in [3], with the aim of defining a much more expressive type system, and a richer set of module operators supporting virtual classes and, more generally, code reuse in presence of unanticipated software modifications.

The key novelties of the approach are the following:

- Type requirements on classes used in a module are *automatically inferred*, without any need for extra type annotations provided by the programmer. We call *smart modules* our modules to stress that, thanks to the type inference algorithm, a basic module (that is, a named collection of class declarations possibly parametric in some classes) can be typechecked and compiled in *full isolation*, without any need that either used classes in source/bytecode form or even a specification file is available. Moreover, these type requirements are *the minimal requirements* on the context needed to safely use a module. This can be achieved by using the recent technology developed in [1], where a type system supporting *principal typings* [32] for Featherweight

Java [22] has been defined. More in detail, in [1] the type information inferred for a class c is a pair consisting of the set of constraints (possibly involving type variables) on all classes used in the code of c (c included) sufficient and necessary for successfully compiling c , and of the type of c which can be easily extracted from its declaration by removing all the method bodies. Furthermore, the notion of *polymorphic bytecode* (bytecode where type annotations can contain type variables) supports separate bytecode generation for classes.

- Classes declared in a module are all implicitly *virtual*, that is, their definition can be later replaced when composing the module with others. The standard way to override components (classes in this case) of mixin modules is by means of a *restrict* operator [10, 4] which allows the deletion of a component definition, in such a way that a new definition can then be provided by merging the module with another. In order to support this mechanism, here type inference uses the same algorithm proposed in [1], but, differently from there (where constraints were incrementally simplified when linking classes together), simplification of constraints is never performed, since they clearly need to be re-checked in case some class declaration is overridden.
- Finally, in addition to classical operators for manipulating mixin modules [10, 4], such as *merge* and *renaming* (already present in [3]), and *restrict*, we introduce here the novel *bind* and *unbind* operators; the former is used to associate module parameters with class declarations already present in the same module, the latter to introduce a new module parameter by removing an existing association. These operators can be expressed in terms of primitive module operators (namely, by using *reduct* [4]), but to our knowledge have never been proposed as high-level constructs in previous languages supporting mixin modules. In this paper, we introduce them (in particular *unbind*) as a powerful means, yet with simple and clean semantics, for supporting unanticipated software evolution, in what it allows an external user to take some code which was designed as closed and to make it parametric, hence available to further composition.

The formalization of the module system is carried out with Featherweight Java as the underlying language. The semantics of module expressions is defined by a set of rewriting rules which allow simplification of each module expression into a basic module. The type system turns out to be sound with respect to the rewriting semantics: the usual subject reduction and progress properties are stated.

The paper is organized as follows. In Section 2 we discuss more in detail the key ideas at the basis of smart modules. Section 3 is a gentle introduction to SMARTJAVAMOD, while Section 4 shows the expressive power of the language and its ability to support software reuse and unanticipated software evolution by using a complete running example. The formalization of the semantics and type system of smart modules for Featherweight Java can be found in Section 5, while Section 6 gives some hints towards a possible implementation. Finally, Section 7

contains pointers to related work and some final remarks and directions for future developments.

2 Smart Modules

Java and other widely used object-oriented languages, like C++, provide only weak module systems. Java packages and C++ namespaces, for instance, are just a means for structuring the namespace, and Java/C++ classes incorporate some, but not all, features found in richer module mechanisms.

For instance, visibility levels are handled at the class level, while a separation of concerns allows a more clean and powerful language design, as already pointed out in [30] (see also [17] for an extended discussion on this point).

Despite the considerable effort invested in studying theoretical foundations and developing new forms of module systems¹, there is relatively little amount of effort, both on the theoretical and implementation side, on applying these ideas to the case of object-oriented languages.

Highly desirable features for a modular approach to software development are:

Separate type-checking A code fragment (module) which needs some services from others should be possibly typed in isolation. In Java the type-checking of a class which relies on other classes requires all the *definitions* of these classes in the environment. For instance, consider the following simple class, representing single-linked lists (where N, not present, is the type of a node):

```
class List {
  N first;
  void addFirst() { first=new N(first); }
}
```

This class cannot be compiled if the definition of class N is unavailable. One reason is the fact that a Java compiler, in order to typecheck the expression `new N(first)`, needs to

- know which (accessible) constructors class N provides,
- choose the most specific among the applicable ones²,
- annotate the generated bytecode for the invocation with the parameter types of the chosen constructor.

For instance, the compiled bytecode might look, from a very abstract point of view, like `new N<<Object>>(first)` if class N provides a constructor which receives an `Object` (and does not provide a more specific constructor).

¹ Let us mention the wide literature about foundations and improvements of Standard ML module system (see, e.g., [25, 21]), the notions of *mixin* (see, e.g., [11, 10, 4]) and *unit* [19] and the type-theoretical analysis of recursion between modules proposed in [16].

² If a most specific constructor cannot be found, then an error is raised and the compilation aborted.

As the name `N` may be seen as a parameter for class `List` and generics³ allow to abstract over types, one might be tempted to transform `List` into `List<N>` in order to compile its definition in isolation:

```
class List<N> {
    N first;
    void addFirst() {
        first=new N(first) ; // ERROR
    }
}
```

Unfortunately, this is not possible for various reasons: first of all, `N` can be used only as a type, that is, instances of `N` cannot be created using `new` operator. Also, inside the body of `List<N>` all you can do with an object of type `N` is determined by its upper bound⁴; this strongly limits the reuse of a module, as we discuss further below. On the contrary, a similar solution works with C++ templates, just because in that case there are no limitations on the use of `N`. However, since typechecking and compilation of templates is postponed at instantiation time, the C++ solution does not meet our modularity requirements.

From an implementation point of view, generics are currently implemented by the Java compiler using type erasure so, all instances of `List<N>` are translated to the same class `List` and, roughly speaking, the identity of `N` is unknown at runtime. Anyway, even if the identity of `N` were not lost at runtime (some extensions of Java, as *LOOJ* [13] and *NextGEN* [14], do not lose it) the fact that in Java constructors are not inherited makes impossible to know which constructors are available given only the upper bound of a type parameter.

Module interfaces A module interface is a specification of the services a given module both needs from and provides to others. As it is well-known, these specifications serve as a formal contract [29] between server and client. Hence, the client can rely on the server specification without any need at looking at the implementation; on the other hand, the correctness of an implementation w.r.t. the specification can be separately checked.

Module expressions A principle which has recently become popular in the programming language community (see, e.g., [27]) is that a module system should have two linguistic levels: a *module language* providing operators for combining software components, constructed on top of a *core language* (following the terminology introduced with Standard ML) for defining module components. The module language should have its own typing rules and be as independent as possible from the core language; even more, it could be in principle instantiated over different core languages.

Module interfaces and separate type-checking usually go hand in hand as the former is used to implement the latter. For instance, *Jiazzi* [28], *Keris* [33] and

³ Introduced in JDK 1.5.

⁴ In this example the upper bound is implicitly the default one: class `Object`.

our JAVAMOD [3] are three proposals which build a module system upon Java, enjoy the desirable features listed before, and require the user to write module interfaces. This last requirement, that is, requiring *the user* to write an explicit module interface is the source of two drawbacks. The more obvious is that the user is forced to write more than she/he would probably like to; the other, less obvious, drawback is the fact that fixing an interface, expressed as a sequence of fields/methods and so on, severely limits the reuse of a module because such an interface, in most cases, is more specific than it needs to be. Let us explain both points by means of an example; consider again the module for defining single-linked lists, sketched before. In JavaMod it can be written as the following basic module:

```
interface ParametricList is
N: { N(N) } ->
    List : { N first ; void addFirst() ; }

module ParametricList is
    class List {
        N first;
        void addFirst() { first=new N(first); }
    }
```

The definition of a basic module is split in two parts: a module interface which describes what a module expects and provides, and a module definition. In this example the module interface describes that, given a (parameter) class `N` which provides a constructor which receives, as the only argument, an object of the same class, the module `ParametricList` provides a class named `List` with an `N` field called `first` and a `void` method called `addFirst`.

Let us consider the following module:

```
interface Node is
-> N : { N(Object) }

module Node is
    class N {
        N(Object) {}
    }
```

When two modules are merged, the expected and provided interfaces are matched for equality, so the module `Node` cannot be merged with module `ParametricList` because class `N` does not provide a constructor which receives an `N`.

Yet, class `List` can be successfully typechecked and compiled in an environment where the only constructor for `N` receives an `Object`; one definition of `N` is as good as the other. So, from the point of view of module `ParametricList`, the most general requirement on the parameter `N` is not simply “to provide a constructor which receives an `N`”, as expressed by a classic module interface, but “to provide a constructor which can be unambiguously called passing an `N`”.

Keeping the module interfaces unaltered and relaxing the check on interface matching is not an option because, as we recalled before, in order to compile the expression `new N(first)`, the compiler must know *exactly* which is the type of the parameter of the chosen constructor. Moreover, relaxing the requirements on interface matching and taking into account overloading resolution would be a very challenging task.

It may appear that trying to make the module interfaces more precise, in order to capture more reuse contexts, would put more burden on the programmer.

Interestingly, we can disburden the programmer from writing module interfaces *and* retain all the desirable features listed before in one fell swoop using inferred type constraints and polymorphic bytecode [1].

As the rest of the paper illustrates, in the new approach we propose here, the programmer does not need to write explicitly the interfaces of the modules a basic module `M` depends on, as the most general requirements needed for a successful compilation of `M` are automatically inferred by the compiler, which can compile `M` in *total isolation*.

3 Language Overview

This section is a brief introduction to smart modules: their main features are presented through some simple, but still meaningful, examples showing their expressive power.

For convenience, the underlying core language used in the examples of this section is Java, but all code could be easily rewritten in C#; however, for simplicity, the formal treatment given in Section 5 will consider only a subset of Java, namely Featherweight Java [22].

3.1 Basic Smart Modules

Let us start our introduction with an example⁵ of declaration of basic module:

```
module LinkedList{ // parametric in N
  class List{
    N first;
    void addFirst(){first=new N(first);}
  }
  class Node{
    N next;
    Node(N n){next=n;}
    N getNext(){return next;}
  }
}
```

A basic module is a collection of class/interface declarations (like `List` and `Node`) possibly parametric in classes/interfaces (like `N`) which can be bound to

⁵ For simplicity, we will keep the examples small and avoid access modifiers.

their definition later. In this example the name `N` is a parameter, as no types named `N` are declared inside the module.

Within this example, the intuition is that `N` *could be* `Node`; indeed, if we replaced all occurrences of `N` with `Node`, then we would obtain the classic example of single-linked lists. Of course, having used a parameter instead of the name of a fixed class, allows us to bind `N` to something more specific than `Node` later, for instance a class `DoubleNode` (which, presumably, extends `Node`).

This particular use of a module parameter roughly corresponds to the idea of type `ThisClass` of LOOJ [13].

Our approach allows a step further though: `N` can be bound to *any class* that provides an accessible constructor which can be invoked passing an object of type `N` as the single argument; that is, a constructor which is declared to receive a `N` or any of its supertypes (superclasses or implemented interfaces, whatever they might be).

Note that even if the programmer does not have to declare:

1. which are the parameters of a module (in this case, `N` is the only one) and
2. what features (methods, fields and so on) `N` is supposed to provide,

our approach allows to compile this basic module in *total isolation* (no other sources or binary `.class` files are needed) and to compose this module with others in a *type safe* manner.

The point is that we do not require the programmer to write the interfaces of used modules because we can *infer* the most general requirements on used modules, which can be later used to check whether module compositions are type safe or not.

For instance, the type inferred for `LinkedList` is

```
{
  List:{cons(N,N,'a)} ->
    extends Object {N first; void addFirst();}
  Node:{exists(N)} ->
    extends Object {N next; Node(N n); N getNext();}
}
```

The type of a module is a mapping from the class names declared in the module to their corresponding class types. In a class type, the set on the left hand side of `->` contains all the type constraints needed for successfully compiling the class, while the right hand side provides all the type information on the class. For instance, the type of `List` means that the class can be successfully compiled if and only if `N` is bound to a class declaration providing a constructor which is the most specific for an argument of type `N` (hence, having a parameter of type `'a`, with `'a` a type variable such that `N <= 'a`). Similar constraints are used for dealing with field access and method invocations.

Note that the type variables which occur in the constraints are all existentially quantified and their scope is limited to the left hand side of `->`; more precisely, the type of `List` should be rewritten as follows:


```
( $\exists$  'a.{cons(N,N,'a)}) ->
  extends Object {N first; void addFirst();}
```

The fact that the right hand side of `->` might contain type variables as well merely depends on the underlying programming language: for instance, only the latest version of Java (JDK 1.5) allows generic classes and polymorphic methods.

Since `N` is used in `Node` as a type, the only constraint is `exists(N)` which requires class `N` to be bound to an existing class declaration.⁶

3.2 Open and Closed Modules

Classes `List` and `Node` inside `LinkedList` cannot be used until the parameter `N` is bound to a class declaration which satisfies the constraints of `List` and `Node`.

```
// the following is a type error,
// since N has not been instantiated yet

LinkedList.List l=new LinkedList.List(); // ERROR
```

A module with parameters, such as `LinkedList`, is called *open*. The module system offers two different operators for instantiating parameters of open modules: *bind* and *merge*.

Bind The bind operator allows self-instantiation. For instance, since class `Node` satisfies all required constraints on `N`, we can define a module `ClosedLinkedList` obtained from `LinkedList` by binding `N` to `Node`:

```
module ClosedLinkedList=bind(LinkedList,{N->Node}) ;
```

Module `ClosedLinkedList` is *closed*, since it has no parameters. Equivalently, `ClosedLinkedList` could have been declared by copying the definition of `LinkedList` and replacing each occurrence of `N` with `Node`.

```
module ClosedLinkedList { // naive cut-and-paste approach

  class List {
    Node first;
    void addFirst() {first=new Node(first);}
  }
  class Node {
    Node next;
    Node(Node n) {next=n;}
    Node getNext() {return next;}
  }
}
```

Now classes `List` and `Node` can be used:

⁶ Note that `cons(N,N,'a)` subsumes `exists(N)` and `N<='a`.

```
ClosedLinkedList.List l=new ClosedLinkedList.List();
```

When closing a module, all type constraints in the class types must be verified, otherwise a type error is issued.⁷ For instance, the expression `bind(LinkedList, {N->List})` is not type correct, since `List` does not satisfy the constraint `cons(List, List, 'a')`.

The reader might expect that the type of a closed module contains no constraints, but this is not the case. Indeed, a closed module is not “sealed” once and for all, but can be reopened using operators *restrict* and *unbind*, which we discuss in Section 3.4. So, for instance, the type of `ClosedLinkedList` is

```
{
List:{cons(Node,Node,'a)} ->
    extends Object {Node first; void addFirst();}
Node:{exists(Node)} ->
    extends Object {Node next; Node(Node n); Node getNext();}
}
```

Finally, note that the fact that a module as `ClosedLinkedList` is closed can be simply deduced from its type: it is immediate to see that all class names contained in the type of `ClosedLinkedList` are declared inside the module itself.

Merge Assume we want to extend the code in `LinkedList` in order to support doubly linked lists. This extension can be isolated in a separate module:

```
module Double { // three parameters: N, List and Node
class DoubleList extends List {
    N last;
    void addLast() {
        N n = new N(last, null) ;
        if (first==null) first = n ;
        if (last!=null) last.next = n ;
        last = n;
    }
    void addFirst() {
        N n=new N(null, first);
        if (first!=null) first.prev = n ;
        first = n ;
        if (last==null) last=n ;
    }
}
class DoubleNode extends Node {
    N prev;
    DoubleNode(N n) {super(n);}
    DoubleNode(N p,N n) {super(n); prev=p;}
    N getPrev() {
        return prev;
    }
}
}
```

⁷ For constraint verification see also Section 3.5.

Now we can define `DoubleLinkedList` by merging `LinkedList` with the extension defined in `Double`:

```
module DoubleLinkedList=merge(LinkedList,Double);
```

Note that the two parameters `List` and `Node` of module `Double` have been automatically instantiated in `DoubleLinkedList` with the corresponding classes declared in `LinkedList`, whereas the parameter `N` is shared (instantiation and sharing of parameters are by class name matching).

Now we can instantiate `N` with `DoubleNode`:

```
module ClosedDoubleLinkedList =
  bind(DoubleLinkedList, {N->DoubleNode}) ;

// this below would be a type error, since
// DoubleList requires N to satisfy
// cons(N,(N,Null),('a','b'))
//
// module ClosedDoubleLinkedList =
//   bind(DoubleLinkedList, {N->Node}) ; // ERROR
```

3.3 Renaming Facilities

Assume that, instead of `N`, `List` and `Node`, the three parameters of module `Double` have been named `DN`, `L` and `N`, respectively. Since, when merging modules, parameter instantiation and sharing is by name matching, we need a renaming operator for correctly defining module `DoubleLinkedList` as above.

```
module DoubleLinkedList =
  merge(LinkedList,
        rename(Double,{N->Node,L->List,DN->N})) ;
```

The *rename* operator allows renaming of a single class name at time, therefore the syntax we have just used is just a convenient shortcut for the more verbose expression:

```
rename(rename(rename(Double,N->Node),L->List),DN->N)
```

This means that renaming of more components is accomplished sequentially from left to right. Both parameters and declared classes can be renamed, however the new name must be unused in order to avoid conflicts; in other words, the operator allows only bijective renaming.

3.4 Virtual Classes and Unanticipated Code Modification

Let us consider again module `ClosedLinkedList` as defined in Section 3.2. As already noted, the type of the module contains constraints on `Node`, even though `Node` is not a parameter. This is because all classes defined in a module are *virtual*, in the sense that their definitions can be overridden; hence, constraints involving a class must be kept in order to be sure that they are still satisfied by

later redefinitions of the other classes. Since the merge operator does not allow class overriding (trying to merge modules which declare classes with the same name is a type error), class redefinition can be accomplished in two steps by means of the *restrict* operator.

For instance, let us assume we want to define a new module `ClosedIntLinkedList` obtained from `ClosedLinkedList` by overriding the definition of `Node`:

```
module IntNode {
  class Node {
    Node next;
    int elem;
    Node(Node n) {next=n;}
    Node(Node n,int e) {next=n;elem=e;}
    Node getNext() {return next;}
    int getElem() {return elem;}
  }
}
module ClosedIntLinkedList =
  merge(restrict(ClosedLinkedList, Node), IntNode);
```

First, we remove the declaration of `Node` in `ClosedLinkedList` by means of the *restrict* operator. Then we add the new declaration contained in `IntNode` with the merge operator.

Now it should be clear why type constraints must always be kept: if we had removed the constraint `cons(Node,Node, 'a)` from the type of `ClosedLinkedList`, we would not be able to correctly typecheck the definition of `ClosedIntLinkedList`.

In general, the *restrict* operator allows the removal of class declarations in a module (trying to remove classes which are not declared is a type error). As for renaming, a convenient shortcut is provided for allowing restriction of multiple classes (however here the order is immaterial).

A preferential merge operator resolving conflicts between class declarations can be obtained as a more powerful form of syntactic shortcut:

```
merge(M1 < M2)
```

is an abbreviation for

```
merge(restrict(M1,{C1,..,Cn}),M2)
```

where `C1,..,Cn` are all the classes declared in both modules.

Finally, *unbind* is another operator which, like *restrict*, enhances code reuse in the presence of unanticipated software modifications. For instance, the class `Node` in module `ClosedLinkedList` as defined in Section 3.2 cannot be directly reused for defining doubly linked nodes as it can be done with the class `Node` in module `LinkedList` declared in Section 3.1. However, the availability of `LinkedList` is not essential, since that module can be obtained from `ClosedLinkedList` by unbinding `Node`.

```
module LinkedList=unbind(ClosedLinkedList,{Node->N})
```

The unbind operator is the inverse of bind: the class to be unbound (`Node` in the example) must be declared in the module while the new name (`N` in the example) must be unused. The effect consists in replacing all non defining occurrences of `Node` with `N`. A defining occurrence is either that immediately following the `class` keyword, or those introducing constructor definitions.

3.5 More on Typing

In this final part more details on type inference for smart modules can be found. We refer to Section 5 and to the related paper [5, 1] for a complete technical treatment of the topic.

Constraint Verification Let us consider the declaration of the following open module:

```
module Vain { // parametric in D
  class C extends D {
    C m() {return new D();}
  }
}
```

In order to be able to use class `C`, the parameter `D` needs to be instantiated. However, it turns out that there is no way to correctly bind `D` to a class declaration: in order to do that the constraint $D \leq C$ in the type of `C` should be satisfied, but class `C` extends `D`, therefore $C < D$ holds which is in contradiction with $D \leq C$. As a matter of fact, module `Vain` is rather useless. From a software engineering point of view, declarations of open modules, like `Vain`, which cannot be closed should be avoided in favor of earlier error detection.

The type inference algorithm defined in [1] is smart enough to reject the declaration of `Vain`; however, it is not so smart to recognize all possible cases. Luckily, verification of constraints is complete in case of closed modules [1]: typing of a closed module succeeds if and only if all its classes are type compatible.

Dot notation As in almost all module systems, smart modules support the dot notation for accessing classes declared in other modules. However two main restrictions apply to the use of the dot notation in order to avoid a too complex type system.

1. As already seen in Section 3.2, dot notation is only allowed for closed modules, in order to avoid run-time errors. For instance, referring to module `LinkedList` defined in Section 3.1, if we try to evaluate the ill-typed expression

```
new LinkedList.List().addFirst()
```

then the exception `NoClassDefFoundError` is thrown when trying to invoke the constructor in the body of `addFirst()`.

2. The most general syntax `ME.c`, where `ME` is a generic module expression, is not allowed; instead, the correct syntactic form is `M.c`, where `M` is a module name. This is the same restriction adopted for the module systems of SML and OCaml in order to avoid a too complex definition of type equivalence. Here we would have the same problems, since in Java a class name denotes a type as well.

Generativity versus Transparency Let us consider the following two module declarations:

```
module ClosedLinkedList =bind(LinkedList,{N->Node})
module ClosedLinkedList2=bind(LinkedList,{N->Node})
```

Should the two types `ClosedLinkedList.List` and `ClosedLinkedList2.List` be considered equivalent? This is a well known issue in ML module systems where the following terminology has been adopted: in case the two types above have to be considered unrelated, the two corresponding module declarations are said to be *generative* (or *opaque*), otherwise they are said to be *transparent*.

In Java all class declarations are generative; for instance, the same class declaration in two different packages always denotes two unrelated types. For analogy one can follow the same approach for modules and, indeed, this is our choice here, but mainly because transparency makes the type system more complex. Let us consider, for instance, the following artificial example:

```
module M1 { // parametric in C
  class A {}
  class B extends C {}
}
module M2=merge({class C{int i;}}, M1);
module M3=merge({class C{boolean b;}}, M1);
```

Let us assume a transparent declaration for `M2` and `M3`. Clearly, `M2.C` and `M3.C` can only be considered unrelated types. Concerning the other two classes `A` and `B`, while `M2.A` and `M3.A` can be safely considered equivalent, the same cannot be done for `M2.B` and `M3.B`, even though `B` comes from the same declaration.

Although the generative approach allows a simpler type system, transparency is advocated in some cases. For this reason, transparent smart modules is an issue that deserves future investigation.

4 Smart Modules at Work

In this section we show the expressive power of smart modules by considering as “benchmark” the classical *expression problem* (or *extensibility problem*). For simplicity we use a slightly more complex variation of Torgersen’s example [31]; we refer to the same paper for a comprehensive treatment of the expression problem which would be out of scope in this paper.

The approach we take here is the classical data-centered one, which is more intuitive and simpler, but also less suitable for adding new methods. Our goal is to show that even with the data-centered approach, smart modules allows addition of new methods with full reuse of the code, at the cost of introducing a parameter (for other approaches on adding methods to existing classes see also [8, 15]).

What follows is an implementation of a type `Exp` of simple integer expressions built on top of literals and addition. The only available methods are `clone()`, which allows cloning of an expression, and `print()`, which displays the expression on the screen.

```
module Exp { // parametric in E

    interface Exp {
        E clone();
        void print();
    }

    class Lit implements Exp {
        int value;
        Lit(int v) {value=v;}
        E clone() {return new Lit(value);}
        void print() {System.out.print(value);}
    }

    class Add implements Exp {
        E left,right;
        Add(E l,E r) {left=l; right=r;}
        E clone() {return new Add(left,right);}
        void print() {
            left.print();
            System.out.print('+');
            right.print();
        }
    }
}
```

The code of the module is very similar to a standard data-centered implementation one would write in a Java package. The only difference is the use of the parameter `E` where the standard implementation would use `Exp` instead. As we will see, this is the only complicity needed for allowing addition of new methods.

The module can be used by a client (for instance a parser) by means of the `bind` operation.

```
module Client { // parametric in Exp, Lit, Add

    class Producer { // typically, the parser
        Exp produce() {return new Add(new Lit(2),new Lit(3));}
    }
}
```

```

}
module Application =
  bind(merge(Exp, Client), {E->Exp}) ;

```

For instance, executing the following well-typed statement

```
new Application.Producer().produce().clone().print();
```

will display 2+3 on the screen, as expected.

Assume now we need to extend our code. Since we have chosen a data-centered approach, adding new methods is more challenging than adding new kinds of expressions, therefore let us assume we want to add a new method `eval()` for evaluating expressions.

We can confine the needed extension into a new module:

```

module Eval{ // parameteric in E, Exp, Lit, Add

  interface EvalExp extends Exp {
    int eval();
  }

  class EvalLit extends Lit implements EvalExp {
    EvalLit(int v) {super(v);}
    int eval() {return value;}
  }

  class EvalAdd extends Add implements EvalExp {
    EvalAdd(E l,E r) {super(l,r);}
    int eval() {return left.eval()+right.eval();}
  }

}

```

A first tentative for instantiating the parameters of `Eval` would consist in directly merging `Eval` with `Exp` and, then, binding `E` to `EvalExp`. However, this is not type correct, since, for instance, class `Lit` in `Exp` requires the constraint `Lit <= E` and this would not hold if we replace `E` with `EvalExp`. The real source of the problem is the fact that now the code of `Lit` and `Add` should refer to classes `EvalLit` and `EvalAdd`. This can be accomplished by unbinding `Lit` and `Add` in module `Exp`.

```

module EvalExp =
  merge(Eval, unbind(Exp, {Lit->EvalLit,Add->EvalAdd}));

module Application2 =
  merge(bind(EvalExp, {E->EvalExp}),
        rename(Client, {Exp->EvalExp, Lit->EvalLit,
                       Add->EvalAdd})) ;

```

Now we can execute the following code:


```

Application2.EvalExp e =
  new Application2.Producer().produce().clone();
e.print();
System.out.println(" = "+e.eval());

```

and, as expected, $2+3 = 5$ will be displayed on the screen.

Finally, note that, because of generativity, the following code is ill-typed:

```

// this below would be a type error:
// Application2.Exp is not comparable with Application.Exp
Application.Exp e2=e; // ERROR

```

5 Smart Modules for Featherweight Java

The syntax of the core language is defined in Fig.1.

$cd ::= \text{class } c \text{ extends } n \{ \text{fds } \text{mds} \}$ ($c \neq \text{Object}$)	class declaration
$\text{fds} ::= n_1 f_1; \dots n_n f_n;$	field declarations
$\text{mds} ::= \text{md}_1 \dots \text{md}_n$	method declarations
$\text{md} ::= \text{mh } \{ \text{return } e; \}$	method declaration
$\text{mh} ::= n_0 m(n_1 x_1, \dots, n_n x_n)$	method header
$e ::= x \mid e.f \mid e_0.m(e_1, \dots, e_n) \mid \text{new } n(e_1, \dots, e_n) \mid (n)e$	expression
$n ::= c \mid M.c$	class name

where field, method and parameter names in fds , mds and mh are distinct

Fig. 1. FJ syntax

We consider the same language as in [1], that is, basically Featherweight Java [22] (FJ in the sequel) hence a functional subset of Java with no primitive types, except that here class constructors are implicitly declared, and class names (references to other classes) appearing in code can be, besides simple names, *qualified names* of the form $M.c$, referring to classes declared in external (closed) modules.

Note that considering a purely functional language here is not a limitation, since modules can only contain classes, and during evaluation of module expressions, classes are not evaluated, hence there would be no side effects even in presence of objects with state.

Every class can contain instance field and method declarations and has only one constructor whose parameters correspond to all class fields (both inherited and declared) in the order of declaration. In class declarations we assume that the name of the class c cannot be `Object`. Method overloading and field hiding are not supported⁸. Expressions are variables, field access, method invocation,

⁸ Just for simplicity and for keeping original FJ definition: it is easy to extend the type system in order to support these features.

instance creation and casting; the keyword `this` is considered a special variable.⁹ Finally, in order to simplify the presentation, we assume field names in `fds`, method names in `mds`, parameter names in `mh` to be distinct.

The syntax of SMARTJAVAMOD is given in Fig.2.

<code>MDS ::= MD₁ ... MD_n</code>	module environment
<code>MD ::= module M is ME</code>	module declaration
<code>ME ::= M cds merge(ME₁, ME₂) restrict(ME, c) rename(ME, c → c') bind(ME, d → c) unbind(ME, c → d)</code>	module expression
<code>cds ::= cd₁ ... cd_n</code>	class declarations
where simple class names in <code>cds</code> are distinct	

Fig. 2. SMARTJAVAMOD syntax

A basic module is a sequence of class declarations, where for simplicity we assume declared class names to be distinct. A class declaration in a basic module can contain qualified class names, that is, references to classes defined in other (closed) modules, and simple class names. Simple class names can be either *defined*, that is, names referring to classes declared in the module, or *deferred*, that is, names for which a declaration is expected to be provided later when composing the module with others. A module with no deferred class names is said to be *closed*. Note that defined class names are not associated with a class declaration in the module once and for all, but their definition can be changed later when composing the module with others. In other words, module components (classes) are all implicitly *virtual*.¹⁰

Module operators include merging two modules, removing a class declaration, renaming a class and two operators which allow to bind a deferred class name to a class declaration inside the module and, conversely, to make a class used inside a module deferred.

Reduction rules for module environments and module expressions (in a given module environment) are those given in Fig.3, plus usual contextual closure for module expressions. We denote by *in*(`cds`) and *out*(`cds`) the deferred and defined class names in class declarations `cds`, respectively.

We call *output* occurrence of a class name `c` in class declarations `cds` only the occurrence in the declaration `class c extends n { fds mds }`, if any; we call *input* occurrence any other occurrence. We denote by `cds[c'/out c]` the class declarations obtained from `cds` by only replacing the output occurrence of `c`, if

⁹ We do not consider the keyword `super` even though used in the examples of Section 3. However, it could be added with no cost at the level of type inference.

¹⁰ We consider only virtual classes in this paper for simplicity, since this is enough for the applications we want to illustrate; however, mixin modules typically support also an operator which allows to make a component *frozen* [10, 4], in such a way that references to this component inside the module can no longer be affected by module composition.

$$\begin{array}{c}
\text{(mdecs)} \frac{\text{MD}_i \rightarrow_{\text{MDS}} \text{MD}'_i}{\text{MD}_1 \dots \text{MD}_n \rightarrow_{\text{MDS}} \text{MD}_1 \dots \text{MD}'_i \dots \text{MD}_n} \text{MDS} = \text{MD}_1 \dots \text{MD}_n \\
\\
\text{(mdec)} \frac{\text{ME} \rightarrow_{\text{MDS}} \text{ME}'}{\text{module M is ME} \rightarrow_{\text{MDS}} \text{module M is ME}'} \\
\\
\text{(mname)} \frac{\text{MDS} = \text{module M}_1 \text{ is ME}_1 \dots \text{module M}_n \text{ is ME}_n}{\text{M} \rightarrow_{\text{MDS}} \text{ME}_i \quad \text{M} = \text{M}_i} \\
\\
\text{(merge)} \frac{}{\text{merge}(\text{cds}_1, \text{cds}_2) \rightarrow_{\text{MDS}} \text{cds}_1 \text{ cds}_2} \text{out}(\text{cds}_1) \cap \text{out}(\text{cds}_2) = \emptyset \\
\\
\text{(restrict)} \frac{}{\text{restrict}(\text{cd}_1 \dots \text{cd}_n, \text{c}) \rightarrow_{\text{MDS}} \text{cd}_1 \dots \text{cd}_{i-1} \text{cd}_{i+1} \dots \text{cd}_n} \text{out}(\text{cd}_i) = \{\text{c}\} \\
\\
\text{(rename)} \frac{\text{c} \in \text{in}(\text{cds}) \cup \text{out}(\text{cds})}{\text{rename}(\text{cds}, \text{c} \rightarrow \text{c}') \rightarrow_{\text{MDS}} \text{cds}[\text{c}'/\text{in c}][\text{c}'/\text{out c}]} \text{c}' \notin \text{in}(\text{cds}) \cup \text{out}(\text{cds}) \\
\\
\text{(bind)} \frac{\text{d} \in \text{in}(\text{cds})}{\text{bind}(\text{cds}, \text{d} \rightarrow \text{c}) \rightarrow_{\text{MDS}} \text{cds}[\text{c}/\text{in d}]} \text{c} \in \text{out}(\text{cds}) \\
\\
\text{(unbind)} \frac{\text{c} \in \text{out}(\text{cds})}{\text{unbind}(\text{cds}, \text{c} \rightarrow \text{d}) \rightarrow_{\text{MDS}} \text{cds}[\text{d}/\text{in c}]} \text{d} \notin \text{in}(\text{cds}) \cup \text{out}(\text{cds})
\end{array}$$

Fig. 3. Reduction rules for module expressions

any, by c' ; we denote by $\text{cds}[c'/in\ c]$ the class declarations obtained from cds by replacing every input occurrence of c by c' . Note that $\text{cds}[c'/out\ c][c'/in\ c]$ corresponds to standard replacement of all occurrences of c by c' in cds , which we will also denote by $\text{cds}[c'/c]$.

Merging two basic modules corresponds to just putting together their class declarations, provided that there are no conflicts (note that deferred class names are shared). The *restrict* operator removes a class declaration, thus making deferred a class name which was defined. The *rename* operator replaces everywhere a class name with a new name. The *bind* operator replaces a deferred class name with the name of a class declared in the module, with the effect of removing a dependency from an external class. Conversely, the *unbind* operator replaces all references to a class declaration inside the module by references to a new name, with the effect of introducing a dependency from an external class.

Types are given in Fig.4.

$\mathcal{M} ::= (M_1, MT_1) \dots (M_n, MT_n)$	module type environment
$MT ::= (\Gamma_1, \delta_1) \dots (\Gamma_n, \delta_n)$	module type
$\Gamma ::= \gamma_1 \dots \gamma_n$	constraints
$\Delta ::= \delta_1 \dots \delta_n$	class type environment
$\delta ::= (c, n, \text{fss}, \text{mss})$	class signature

Fig. 4. Types

A *module type environment* is a sequence of *module type assignments*, which are pairs consisting of a module name and a module type. A *module type* is the type information needed to safely use a module in a context, and consists of a sequence of *class type assignments*. A class type assignment is the type information on a single class declared in the module, and consists in a sequence of *constraints* and a *class signature*.

Constraints γ [1] express expectations on other classes used inside the class, and are formally defined and explained in the Appendix. Constraints may contain *type variables* α which denote arbitrary class names. For instance, $\phi(c, f, \alpha)$ means that class c is expected to provide a field named f of an arbitrary type. If the class for which this constraint has been inferred is then combined with a class named c providing a field named f of type, say, d , then the constraint is satisfied and the type variable α is instantiated to d .

A class type environment is a sequence of class signatures. Class signatures are the type information which can be extracted from a class declaration and consist of a simple class name (the name of the declared class), a class name (the name of the parent class), a sequence of *field signatures* (type and name of declared fields) and a sequence of *method signatures* (return type, name and parameter types of declared methods). The formal definition is provided in Fig.8 in the Appendix.

Typing rules are given in Fig.5.

$$\begin{array}{c}
\text{(mdec)} \frac{\mathcal{M} \vdash \text{MD}_i : (\text{M}_i, \text{MT}_i) \forall i \in 1..n}{\vdash \text{MDS} : \mathcal{M}} \quad \begin{array}{l} \text{MDS} = \text{MD}_1 \dots \text{MD}_n \\ \mathcal{M} = (\text{M}_1, \text{MT}_1) \dots (\text{M}_n, \text{MT}_n) \\ \rightarrow_{\text{MDS}}^+ \text{acyclic} \end{array} \\
\\
\text{(mdec)} \frac{\mathcal{M} \vdash \text{ME} : \text{MT}}{\mathcal{M} \vdash \text{module M is ME} : (\text{M}, \text{MT})} \\
\\
\text{(mname)} \frac{\mathcal{M} = (\text{M}_1, \text{MT}_1) \dots (\text{M}_n, \text{MT}_n)}{\mathcal{M} \vdash \text{M} : \text{MT}_i \quad \text{M}_i = \text{M}} \\
\\
\text{(basic)} \frac{\begin{array}{l} \vdash \text{cd}_i : \Gamma_i \mid \delta_i \forall i \in 1..n \\ \Delta \vdash \Gamma_i \rightsquigarrow \Gamma'_i \forall i \in 1..n \\ \Delta \delta_1 \dots \delta_n \vdash \Gamma'_1 \dots \Gamma'_n \diamond \end{array}}{\mathcal{M} \vdash \text{cd}_i^{i \in 1..n} : (\Gamma'_i, \delta_i)_{i \in 1..n}} \quad \Delta = \text{ctenv}(\mathcal{M}) \\
\\
\text{(merge)} \frac{\begin{array}{l} \mathcal{M} \vdash \text{ME}_1 : \text{MT}_1 \\ \mathcal{M} \vdash \text{ME}_2 : \text{MT}_2 \\ \Delta \Delta_1 \Delta_2 \vdash \Gamma \diamond \end{array}}{\mathcal{M} \vdash \text{merge}(\text{ME}_1, \text{ME}_2) : \text{MT}_1 \text{ MT}_2} \quad \begin{array}{l} \Delta = \text{ctenv}(\mathcal{M}) \\ \Delta_1 = \text{ctenv}(\text{MT}_1) \\ \Delta_2 = \text{ctenv}(\text{MT}_2) \\ \Gamma = \text{cnstrs}(\text{MT}_1 \text{ MT}_2) \\ \text{out}(\text{MT}_1) \cap \text{out}(\text{MT}_2) = \emptyset \end{array} \\
\\
\text{(restrict)} \frac{\mathcal{M} \vdash \text{ME} : (\Gamma_1, \delta_1) \dots (\Gamma_n, \delta_n)}{\mathcal{M} \vdash \text{restrict}(\text{ME}, \text{c}) : (\Gamma_1, \delta_1) \dots (\Gamma_{i-1}, \delta_{i-1}) (\Gamma_{i+1}, \delta_{i+1}) (\Gamma_n, \delta_n)} \quad \text{out}(\delta_i) = \text{c} \\
\\
\text{(rename)} \frac{\mathcal{M} \vdash \text{ME} : \text{MT}}{\mathcal{M} \vdash \text{rename}(\text{ME}, \text{c} \rightarrow \text{c}') : \text{MT}[\text{c}'/\text{in c}][\text{c}'/\text{out c}]} \quad \begin{array}{l} \text{c} \in \text{in}(\text{c ds}) \cup \text{out}(\text{c ds}) \\ \text{c}' \notin \text{in}(\mathcal{M}) \cup \text{out}(\mathcal{M}) \end{array} \\
\\
\text{(bind)} \frac{\begin{array}{l} \mathcal{M} \vdash \text{ME} : \text{MT} \\ \Delta \Delta' \vdash \Gamma \diamond \end{array}}{\mathcal{M} \vdash \text{bind}(\text{ME}, \text{d} \rightarrow \text{c}) : \text{MT}[\text{c}/\text{in d}]} \quad \begin{array}{l} \Delta = \text{ctenv}(\mathcal{M}) \\ \Delta' = \text{ctenv}(\text{MT}[\text{c}/\text{in d}]) \\ \Gamma = \text{cnstrs}(\text{MT}[\text{c}/\text{in d}]) \\ \text{d} \in \text{in}(\text{MT}) \\ \text{c} \in \text{out}(\text{MT}) \end{array} \\
\\
\text{(unbind)} \frac{\mathcal{M} \vdash \text{ME} : \text{MT}}{\mathcal{M} \vdash \text{unbind}(\text{ME}, \text{c} \rightarrow \text{d}) : \text{MT}[\text{d}/\text{in c}]} \quad \begin{array}{l} \text{c} \in \text{out}(\text{MT}) \\ \text{d} \notin \text{in}(\text{MT}) \cup \text{out}(\text{MT}) \end{array}
\end{array}$$

Fig. 5. Typing rules

In rule (mdecs), a sequence of module declarations is well-formed and produces a given module type environment if under this module type environment each module declaration is well-formed and has the assigned module type. Note that this metarule is recursive to allow mutual references to module components, as in the following example:

```
module M1 is class C { M2.D f; ... }
module M2 is class D { M1.C g; ... }
```

However, mutual references to module names, as in

```
module M1 is rename C by D in M2
module M2 is rename D by C in M1
```

obviously make no sense, hence must be rejected by the type system. To this end, the side condition in rule (mdecs) requires the relation $\rightarrow_{\text{MDS}}^+$ to be acyclic, where $M_1 \rightarrow_{\text{MDS}} M_2$ if there is in MDS a declaration `module M1 is ME` with M_2 subterm of ME, and $\rightarrow_{\text{MDS}}^+$ denotes the transitive closure of \rightarrow_{MDS} .

Rule (mdec) and (mname) are straightforward. We denote by $mname(\text{MT})$ the first component of a module type (the module name).

In rule (basic), the type of a basic module (sequence of class declarations) can be inferred by the following steps:

- Each class declaration is separately type-checked, obtaining the constraints on other classes and the class signature. The formal definition of the judgment $\vdash \text{cd} : \Gamma \mid \delta$, introduced in [1], is reported in Fig.8 in the Appendix.
- For each declared class, it is checked that constraints on qualified class names, that is, class names of the form `M.c`, are satisfied by components of other modules, using their type information provided by the module type environment \mathcal{M} , and these constraints are removed. Formally, constraints Γ_i are simplified to Γ'_i under the class type environment Δ which can be extracted from \mathcal{M} . Note that a class signature for `M.c` can be extracted from the module type of `M` in \mathcal{M} only if `M` has a `c` component and, moreover, is closed. This is expressed by the formal definition of $ctenv$ in Fig.6, where $in(\text{MT})$ and $out(\text{MT})$ denotes the deferred and defined class names appearing in MT, respectively (defined class names are those appearing as first components of class signatures), and $\delta[c'/c]$ denotes replacement of all occurrences of `c` by `c'` in δ . For the formal definition of the judgment $\Delta \vdash \Gamma \rightsquigarrow \Gamma'$ we refer to [1].
- Then, it is checked that mutual constraints are satisfied by the class declarations in the module. Formally, constraints $\Gamma'_1 \dots \Gamma'_n$ left from the previous step are checked under the class type environment consisting of Δ and the class signatures extracted from the class declarations. Note that these constraints are checked but not simplified, since, as already explained, they need to be checked again in case a class declaration will be later overridden. The notation $\Delta \vdash \Gamma \diamond$ is an abbreviation for $\Delta \vdash \Gamma \rightsquigarrow \Gamma'$ for some Γ' .

$$\begin{aligned}
ctenv((M_1, MT_1) \dots (M_n, MT_n)) &= ctenv((M_1, MT_1)) \dots ctenv((M_n, MT_n)) \\
ctenv((M, MT)) &= ctenv(MT)[M.c_i/c_i^{i \in 1..n}] \text{ where } \{c_1, \dots, c_n\} = out(MT) \text{ if } in(MT) = \emptyset, \\
ctenv((M, MT)) &= \Lambda \text{ if } in(MT) \neq \emptyset \\
ctenv((\Gamma_1, \delta_1) \dots (\Gamma_n, \delta_n)) &= \delta_1 \dots \delta_n \\
cnstrs((\Gamma_1, \delta_1) \dots (\Gamma_n, \delta_n)) &= \Gamma_1 \dots \Gamma_n
\end{aligned}$$

Fig. 6. Functions *ctenv* and *cnstrs*

In rule (merge), the operator can be safely applied only if the arguments have no conflicting class declarations, and, moreover, mutual constraints are satisfied by the class declarations in the two modules. Formally, constraints required by either argument are checked under the class type environment consisting of Δ and the class signatures extracted from both module types. The resulting module type is simply obtained by putting together the module types of the arguments. As in (basic), constraints are checked but not simplified.

In rule (restrict), the operator can be safely applied only if the class to be removed actually is a defined class of the module. The resulting module type is obtained by removing the corresponding class type assignment from the module type of the argument.

In rule (rename), the operator can be safely applied only if the class to be renamed is either a deferred or defined class of the module, and the new name is unused. The resulting module type is obtained by correspondingly renaming the module type of the argument.

In rule (bind), the operator can be safely applied only if a deferred class name d is bound to a defined class name c , and, moreover, constraints which previously involved d are satisfied by c . Formally, constraints extracted from the module type, where d has been replaced by c , are checked under the class type environment consisting of Δ and the class signatures extracted from the module type, where also d has been replaced by c . The resulting module type is obtained by replacing d by c in the module type of the argument. As in (basic) and (merge), constraints are checked but not simplified.

In rule (unbind), the operator can be safely applied only if a defined class name c is unbound, by replacing all its input occurrences by a fresh name d . The resulting module type is obtained by analogously replacing all input occurrences of c by d in the module type of the argument (in analogy to what we have defined for class declarations, in a module type output occurrences are only those appearing as first components of class signatures).

In summary, the only operators which require a check of constraints are basic module, merge and bind, since these are the cases in which class names which were deferred (or just taken in isolation in the case of basic module) can become defined, hence some new type information on them can be provided which could

possibly contradict some required constraints. Restrict corresponds to remove some constraints, rename and unbind just to a syntactic manipulation of them.

Also note that constraint checking in the above rules is sometimes redundant, due to the fact that, as already explained, constraints are checked but not simplified in rules (basic), (merge), and (bind). An alternative type system could also keep trace, for each module, of the simplified constraints (that is, those which cannot be checked inside the module itself, since they still involve deferred classes). This would allow a more efficient approach in these rules: for instance, in rule (merge) it would be enough to check the simplified constraints of both modules, whereas in the rule presented here also constraints expressing internal consistency of both modules are checked again; however, with this approach it would be necessary to calculate again simplified constraints in rules (restrict) and (unbind), whereas no constraint simplification step is required for these rules in the current approach.

The relevance of the type system presented until now is that it supports *compositional compilation* of modules. This means that it is possible for the programmer to write and compile a SMARTJAVAMOD module in isolation¹¹, and then to combine it with other modules by just checking that mutual assumptions are satisfied, without any need of reinspecting code. In this sense, SMARTJAVAMOD is as an expressive module language layer constructed on top of the type system in [1]. Here only compositional compilation of single classes and *fragments*, consisting of either a single class or, inductively, a concatenation of fragments (as a matter of fact, an implicit *merge* operator), was considered.

Of course, compositional compilation can be safely used only in place of *global* compilation (that is, compilation of a sequence of classes in a context where all used classes are available in either source or binary form, as standard Java compilers do) only if module typing rules guarantee that, whenever by module composition we obtain in the end a closed module expression, then this module expression reduces to a well-typed self-contained sequence of classes. This is formally expressed, in a more general formulation which also takes into account module environment, by Theorem 3 below, which, as usually, can be proved by means of subject reduction and progress properties stated in Theorem 1 and 2, respectively.

Theorem 1 (Subject reduction).

- If $\vdash MDS : \mathcal{M}$, $MDS \rightarrow MDS'$, then $\vdash MDS' : \mathcal{M}$.
- If $\vdash MDS : \mathcal{M}$, $\mathcal{M} \vdash ME : MT$, $ME \rightarrow_{MDS} ME'$, then $\mathcal{M} \vdash ME' : MT$.

We say that a module declaration `module M is ME` is a *basic module declaration* if `ME` is a basic module; we say that a module environment $MD_1 \dots MD_n$ is a *basic module environment* if each MD_i is a basic module declaration. For MDS basic module environment, we denote by $cenv(MDS)$ the sequence of class declarations consisting of all the class declarations in some closed module, say `M`, where simple

¹¹ That is, in a context where classes used inside the module are not available; however, note that used *modules* must be available.

class names have been qualified by M , as formally defined in Fig.7. Analogously, $mt(\mathcal{M})$ denotes the sequence of class type assignments extracted from a module type environment.

$$\begin{aligned}
cenv(\text{MD}_1 \dots \text{MD}_s) &= cenv(\text{MD}_1) \dots cenv(\text{MD}_s) \\
cenv(\text{module } M \text{ is } cds) &= cds[M.c_i/c_i^{i \in 1..n}] \text{ where } \{c_1, \dots, c_n\} = out(cds) \text{ if } in(cds) = \emptyset \\
cenv(\text{module } M \text{ is } cds) &= \Lambda \text{ if } in(cds) \neq \emptyset \\
mt((M_1, MT_1) \dots (M_n, MT_n)) &= mt(M_1, MT_1) \dots mt(M_n, MT_n) \\
mt(M, MT) &= MT[M.c_i/c_i^{i \in 1..n}] \text{ where } \{c_1, \dots, c_n\} = out(MT) \text{ if } in(MT) = \emptyset \\
mt(M, MT) &= \Lambda \text{ if } in(MT) \neq \emptyset
\end{aligned}$$

Fig. 7. Functions $cenv$ and mt

Theorem 2 (Progress).

- If $\vdash MDS : \mathcal{M}$, MDS not a basic module environment, then $MDS \rightarrow MDS'$ for some MDS' .
- If $\mathcal{M} \vdash ME : MT$, ME not a basic module, then $ME \rightarrow_{MDS} ME'$ for some ME' .

Theorem 3. If $\vdash MDS : \mathcal{M}$, $\mathcal{M} \vdash ME : MT$, $in(MT) = \emptyset$, then $MDS \rightarrow^* MDS'$ with MDS basic module environment, $ME \rightarrow_{MDS}^* cds$, $\Lambda \vdash cenv(MDS') cds : mt(\mathcal{M}) MT$ with $in(mt(\mathcal{M}) MT) = \emptyset$.

Theorem 3 states that, starting from a well-formed closed module expression in a well-formed module environment, we can always obtain a well-formed closed sequence of class declarations, consisting, roughly speaking, of the class declarations we get by reducing the module expression, together with those obtained by “flattening” class declarations in closed modules in the environment. Hence, the reduction semantics of an FJ expression in the context of a closed module expression ME and a module environment MDS can be defined as the reduction semantics of the expression in the context of the FJ program obtained by reducing ME and reducing and flattening MDS , and the result above guarantees type safety.

6 Implementation issues

In this section we discuss how a prototype compiler for SMARTJAVAMOD could be implemented.

One of the key points of our approach is the ability of compiling Java classes and interfaces¹² in total isolation, without any burden on the programmers’ side. The idea, as we already said, is to generate polymorphic bytecode, along with

¹² Not to be confused with module interfaces.

the type constraints representing the inferred requirements. We refer to [1] for a more detailed discussion of this topic.

A basic module is a collection of one or more classes/interfaces, which can be defined in one single file, as the examples seen so far may have suggested, or in many files, using one module definition file which “points to” standard Java sources (this feature comes handy to include plain Java classes in SMARTJAVAMOD modules).

A non basic module is defined in a single file which contains the association between the module name and its defining module expression.

When compiled, all kinds of modules share the same file format. The binary form of compiled modules is very similar to Jiazzi’s one [28]; that is, a compiled module is an archive (a JAR file) which contains:

- the binary form of all classes/interfaces contained in the module, and
- the type constraints needed to ensure a type-safe linking.

If a module is open, then binary classes consist of polymorphic bytecode which cannot be directly loaded, much less executed, by a standard JVM (Java Virtual Machine). This is not a problem though, as an open module is incomplete and could not be run anyway.

When modules are composed by module operators, the type constraints are checked for consistency and if no errors are detected, a new module is created. Unfortunately, as said in Section 3.5, some errors could go undetected as long as modules remains open. Luckily, verification of constraints is complete in case of closed modules [1]. Enhancing the verification algorithm in order to capture inconsistencies as soon as possible is a challenging task which deserves future work.

In a closed module binary classes are stored in plain standard bytecode, so a binary module can be loaded and executed on any JVM. Moreover, module names are mapped into Java package names so, for instance, a class `C` contained in a module called `m` can be referenced by any Java class using the name `m.C`. This makes possible to use a compiled SMARTJAVAMOD (closed) module with any Java compiler (because type constraints are simply ignored by standard tools).

A feature, which could be beneficial for a real-world implementation, is the ability to add, along inferred type constraints, some constraints explicitly specified by the programmer. While this feature is unnecessary from a strictly theoretical point of view, as the most general requirements are always inferred, it could come handy when programmers needs to enforce particular module interfaces in order to interact with legacy code. Fortunately, this feature is very easy to add, as the constraint checking algorithm does not care whether a constraint has been inferred or written by the programmer. So, the compiler just needs to add the specified constraints to inferred ones to implement this feature.

7 Conclusion

We have presented SMARTJAVAMOD, a language of mixin modules constructed on top of Java which supports compositional compilation [1]. The semantics is given by reduction into pure Java, and the type system guarantees that the collection of classes obtained in the end is a well-formed Java program.

We have obtained a module language which is extremely flexible and allows to express a variety of constructs supporting software reuse and extensibility.

There are, to our knowledge, few attempts at integrating a true module system with Java. They include the simple extension of the package system proposed in [7], Jiazzi [28], Keris [33] and our own work in [3]. Instead there are many proposals for adding expressiveness to Java-like languages by extending the language in various ways, e.g., by introducing mixin classes [2], and virtual classes [12]. On the other hand, there exist several languages supporting modules in the context of object-oriented programming with structured types [26, 18]; however, nominal types, as in Java, introduce different problems.

Concerning type inference for Java/Javascript-like languages, we mention some recent work [6], which, however, considers a very different setting (object based with structural subtyping).

The approach presented here introduces the following key novelties. First, *compositional compilation* allows to typecheck a module in isolation, inferring the type constraints needed for safely using it in any context, without any need for the user to specify an interface. Secondly, the module language supports *mixin modules*, that is, module operators which allow to redefine components. Finally, this module layer allows to express a variety of constructs, including mixin classes and virtual classes, without changing the core level (that is, keeping pure Java as core language).

There are two main directions for further research we want to explore. On the one hand, we plan to investigate in a more formal way how derived constructs, such as mixin classes, virtual classes and traits, can be encoded in SMARTJAVAMOD.

On the other hand, much remains to be investigated on the side of the implementation, in particular if one wants (as it seems desirable) to obtain a running extension actually constructed on top of Java. In particular, as already said in the paper, the verification of constraints is complete only in case of closed modules, so enhancing the algorithm in order to capture inconsistencies as soon as possible deserves future work.

Furthermore, a drawback of pure constraint inference is that there is no way to discover trivial typing errors, such as invoking the wrong method name, before linking since any reference to a missing component is just interpreted as a constraint. To avoid this, constraint inference should be in practice complemented by additional instruments to improve the practicality of the approach. For instance, compilers could print partial information extracted from the inferred constraints (which are heavy in their complete form) that can be checked by programmers, who should also be allowed to explicitly add constraints to the inferred ones.

References

1. D. Ancona, F. Damiani, S. Drossopoulou, and E. Zucca. Polymorphic bytecode: Compositional compilation for Java-like languages. In *ACM Symp. on Principles of Programming Languages 2005*, January 2005.
2. D. Ancona, G. Lagorio, and E. Zucca. Jam—designing a Java extension with mixins. *ACM Transactions on Programming Languages and Systems*, 25(5):641–712, September 2003.
3. D. Ancona and E. Zucca. True modules for Java-like languages. In J. Knudsen, editor, *ECOOP'01 - European Conference on Object-Oriented Programming*, number 2072 in Lecture Notes in Computer Science, pages 354–380. Springer, 2001.
4. D. Ancona and E. Zucca. A calculus of module systems. *Journ. of Functional Programming*, 12(2):91–132, 2002.
5. D. Ancona and E. Zucca. Principal typings for Java-like languages. In *ACM Symp. on Principles of Programming Languages 2004*, pages 306–317. ACM Press, January 2004.
6. C. Anderson, P. Giannini, and S. Drossopoulou. Towards type inference for Javascript. In *ECOOP'05 - Object-Oriented Programming*, Lecture Notes in Computer Science. Springer, 2005. To appear.
7. L. Bauer, A. Appel, and E. Felten. Mechanisms for secure modular programming in Java. Technical Report CS-TR-603-99, Department of Computer Science, Princeton University, July 1999.
8. A. Bergel, S. Ducasse, O. Nierstrasz, and R. Wuyts. Classboxes: Controlling visibility of class extensions. *Computer Languages, Systems and Structures*, Sept. 2005.
9. V. Bono, A. Patel, and V. Shmatikov. A core calculus of classes and mixins. In R. Guerraoui, editor, *ECOOP'99 - European Conference on Object-Oriented Programming*, number 1628 in Lecture Notes in Computer Science, pages 43–66. Springer, 1999.
10. G. Bracha. *The Programming Language JIGSAW: Mixins, Modularity and Multiple Inheritance*. PhD thesis, Department of Comp. Sci., Univ. of Utah, 1992.
11. G. Bracha and W. Cook. Mixin-based inheritance. In *ACM Symp. on Object-Oriented Programming: Systems, Languages and Applications 1990*, volume 25(10) of *SIGPLAN Notices*, pages 303–311. ACM Press, October 1990.
12. K. Bruce, M. Odersky, and P. Wadler. A statically safe alternative to virtual types. In *ECOOP'98 - European Conference on Object-Oriented Programming*, number 1445 in Lecture Notes in Computer Science, pages 523–549, 1998.
13. K. B. Bruce and J. N. Foster. LOOJ: Weaving LOOM into Java. In *ECOOP'04 - Object-Oriented Programming*, number 3086 in Lecture Notes in Computer Science, pages 389–413, 2004.
14. R. Cartwright and G. L. Steele, Jr. Compatible genericity with run-time types for the Java programming language. In C. Chambers, editor, *ACM Symp. on Object-Oriented Programming: Systems, Languages and Applications 1998*, pages 201–215. ACM Press, 1998.
15. C. Clifton, T. Millstein, G. T. Leavens, and C. Chambers. MultiJava: Design rationale, compiler implementation, and applications. Technical Report 04-01b, Iowa State University, Dept. of Computer Science, Dec. 2004. Accepted for publication, pending revision.
16. K. Cray, R. Harper, and S. Puri. What is a recursive module? In *PLDI'99 - ACM Conf. on Programming Language Design and Implementation*, 1999.

17. K. Fisher and J. Reppy. The design of a class mechanism for MOBY. In *PLDI'99 - ACM Conf. on Programming Language Design and Implementation*, pages 37–49, 1999.
18. K. Fisher and J. Reppy. Extending MOBY with inheritance-based subtyping. In *ESOP 2000 - European Symposium on Programming 2000*, number 1850 in Lecture Notes in Computer Science, pages 83–107. Springer, 2000.
19. M. Flatt and M. Felleisen. Units: Cool modules for HOT languages. In *PLDI'98 - ACM Conf. on Programming Language Design and Implementation*, pages 236–248, 1998.
20. M. Flatt, S. Krishnamurthi, and M. Felleisen. Classes and mixins. In *ACM Symp. on Principles of Programming Languages 1998*, pages 171–183. ACM Press, 1998.
21. R. Harper and M. Lillibridge. A type theoretic approach to higher-order modules with sharing. In *ACM Symp. on Principles of Programming Languages 1994*, pages 127–137. ACM Press, 1994.
22. A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *ACM Symp. on Object-Oriented Programming: Systems, Languages and Applications 1999*, pages 132–146, November 1999.
23. G. Kiczales. Aspectj(tm): Aspect-oriented programming in Java. In M. Aksit, M. Mezini, and R. Unland, editors, *NetObjectDays 2002*, number 2591 in Lecture Notes in Computer Science. Springer, 2003.
24. G. Kiczales, J. Lamping, A. Mendhekar, C. Maeda, C. V. Lopes, J.-M. Loingtier, and J. Irwin. Aspect-oriented programming. In M. Aksit and S. Matsuoka, editors, *ECOOP*, volume 1241 of *Lecture Notes in Computer Science*, pages 220–242. Springer, 1997.
25. X. Leroy. Manifest types, modules and separate compilation. In *ACM Symp. on Principles of Programming Languages 1994*, pages 109–122. ACM Press, 1994.
26. X. Leroy. The Objective Caml system (release 2.00). Available at <http://paulliac.inria.fr/caml>, August 1998.
27. X. Leroy. A modular module system. *Journal of Functional Programming*, 10(3):269–303, May 2000.
28. S. McDirmid, M. Flatt, and W. Hsieh. Jiazzi: New age components for old fashioned Java. In *ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA 2001)*. ACM Press, October 2001.
29. B. Meyer. *Eiffel: The language*. Prentice Hall, 1992.
30. C. Szyperski. Import is not inheritance. why we need both: Modules and classes. In O. L. Madsen, editor, *ECOOP'92 - European Conference on Object-Oriented Programming*, number 615 in Lecture Notes in Computer Science, pages 19–32. Springer, 1992.
31. M. Torgersen. The expression problem revisited. In M. Odersky, editor, *ECOOP'04 - Object-Oriented Programming*, number 3086 in Lecture Notes in Computer Science, pages 123–143. Springer, 2004.
32. J. Wells. The essence of principal typings. In *International Colloquium on Automata, Languages and Programming 2002*, number 2380 in Lecture Notes in Computer Science, pages 913–925. Springer, 2002.
33. M. Zenger. Evolving software with extensible modules. In *International Workshop on Unanticipated Software Evolution*, Malaga, Spain, June 2002.

Typing rules for classes [1] are formally defined in Fig.8. Constraints have the following informal meaning:

- $t \leq t'$ means “ t is a subtype of t' ”.

- $\exists \mathbf{n}$ means “ \mathbf{n} is defined”.
- $\phi(\mathbf{t}, \mathbf{f}, \mathbf{t}')$ means “ \mathbf{t} provides field \mathbf{f} with type \mathbf{t}' ”.
- $\mu(\mathbf{t}, \mathbf{m}, \bar{\mathbf{t}}, (\mathbf{t}', \bar{\mathbf{t}}'))$ means “ \mathbf{t} provides method \mathbf{m} applicable to arguments of type $\bar{\mathbf{t}}$, with return type \mathbf{t}' and parameters of type $\bar{\mathbf{t}}'$ ”.
- $\kappa(\mathbf{n}, \bar{\mathbf{t}}, \bar{\mathbf{t}}')$ means “ \mathbf{n} provides constructor applicable to arguments of type $\bar{\mathbf{t}}$, with parameters of type $\bar{\mathbf{t}}'$ ”.
- $\mathbf{n} \sim \mathbf{t}$ means “ \mathbf{n} and \mathbf{t} are comparable” (this constraint is generated when compiling a cast).

Note that both the constraints $\mu(\mathbf{t}, \mathbf{m}, \bar{\mathbf{t}}, (\mathbf{t}', \bar{\mathbf{t}}'))$ and $\kappa(\mathbf{n}, \bar{\mathbf{t}}, \bar{\mathbf{t}}')$ implicitly include the constraint $\bar{\mathbf{t}} \leq \bar{\mathbf{t}}'$.

We write $type(\mathbf{fds})$ and $type(\mathbf{mds})$ to denote the set of field signatures and the set of method signatures extracted from the field declarations \mathbf{fds} and from the method declarations \mathbf{mds} , respectively. The straightforward definition of $type$ has been omitted.

$\gamma ::= \mathbf{t} \leq \mathbf{t}' \mid \exists \mathbf{n} \mid \phi(\mathbf{t}, \mathbf{f}, \mathbf{t}') \mid \mu(\mathbf{t}, \mathbf{m}, \bar{\mathbf{t}}, (\mathbf{t}', \bar{\mathbf{t}}')) \mid \kappa(\mathbf{n}, \bar{\mathbf{t}}, \bar{\mathbf{t}}') \mid \mathbf{n} \sim \mathbf{t}$ constraint	
$\mathbf{t} ::= \mathbf{n} \mid \alpha$	expression type
$\bar{\mathbf{t}} ::= \mathbf{t}_1 \dots \mathbf{t}_n$	expression types
$\mathbf{fss} ::= \mathbf{fs}_1 \dots \mathbf{fs}_n$	field signatures
$\mathbf{fs} ::= \mathbf{n} \mathbf{f}$	field signature
$\mathbf{mss} ::= \mathbf{ms}_1 \dots \mathbf{ms}_n$	method signatures
$\mathbf{ms} ::= \mathbf{n} \mathbf{m}(\bar{\mathbf{n}})$	method signature
$\bar{\mathbf{n}} ::= \mathbf{n}_1 \dots \mathbf{n}_n$	class names
$\bar{\alpha} ::= \alpha_1 \dots \alpha_n$	type variables
$\Pi ::= (\mathbf{x}_1, \mathbf{n}_1) \dots (\mathbf{x}_n, \mathbf{n}_n)$	local environment

$(class) \frac{\vdash \mathbf{fds} : \Gamma \quad \mathbf{c} \vdash \mathbf{mds} : \Gamma'}{\vdash \mathbf{class} \ \mathbf{c} \ \mathbf{extends} \ \mathbf{n} \ \{ \ \mathbf{fds} \ \mathbf{mds} \} : \Gamma, \Gamma', \exists \mathbf{n} \mid (\mathbf{c}, \mathbf{n}, \mathbf{fss}, \mathbf{mss})}$	$type(\mathbf{mds}) = \mathbf{mss}$ $type(\mathbf{fds}) = \mathbf{fss}$
$(fields) \frac{\vdash \mathbf{fd}_i : \Gamma_i \quad \forall i \in 1..n \quad n \neq 1}{\vdash \mathbf{fd}_1 \dots \mathbf{fd}_n : \Gamma_1 \dots \Gamma_n}$	$(field) \frac{}{\vdash \mathbf{n} \ \mathbf{f}; : \exists \mathbf{n}}$
$(methods) \frac{\mathbf{c} \vdash \mathbf{md}_i : \Gamma_i \quad \forall i \in 1..n \quad n \neq 1}{\mathbf{c} \vdash \mathbf{md}_1 \dots \mathbf{md}_n : \Gamma_1 \dots \Gamma_n}$	
$(method) \frac{\mathbf{x}_1:\mathbf{n}_1 \dots \mathbf{x}_n:\mathbf{n}_n, \mathbf{this}:\mathbf{c} \vdash \mathbf{e} : \mathbf{t} \mid \Gamma}{\mathbf{c} \vdash \mathbf{n}_0 \ \mathbf{m}(\mathbf{n}_1 \ \mathbf{x}_1 \dots \mathbf{n}_n \ \mathbf{x}_n) \ \{ \ \mathbf{return} \ \mathbf{e}; \} : \Gamma, \mathbf{t} \leq \mathbf{n}_0, \exists \mathbf{n}_i^{i \in 0..n}}$	
$(parameter) \frac{}{\Pi \vdash \mathbf{x} : \mathbf{n}_i \mid \Delta \quad \mathbf{x} = \mathbf{x}_i}$	$(field \ access) \frac{\Pi \vdash \mathbf{e} : \mathbf{t} \mid \Gamma}{\Pi \vdash \mathbf{e}.\mathbf{f} : \alpha \mid \Gamma, \phi(\mathbf{t}, \mathbf{f}, \alpha)} \ \alpha \ \text{fresh}$
$(meth \ call) \frac{\Pi \vdash \mathbf{e}_0 : \mathbf{t}_0 \mid \Gamma_0 \quad \Pi \vdash \mathbf{e}_i : \mathbf{t}_i \mid \Gamma_i \quad \forall i \in 1..n}{\Pi \vdash \mathbf{e}_0.\mathbf{m}(\mathbf{e}_1 \dots \mathbf{e}_n) : \beta \mid \Gamma_0 \Gamma_1 \dots \Gamma_n, \mu(\mathbf{t}_0, \mathbf{m}, \mathbf{t}_1 \dots \mathbf{t}_n, (\beta, \bar{\alpha}))}$	$\beta, \bar{\alpha} \ \text{fresh}$
$(new) \frac{\Pi \vdash \mathbf{e}_i : \mathbf{t}_i \mid \Gamma_i \quad \forall i \in 1..n}{\Pi \vdash \mathbf{new} \ \mathbf{n}(\mathbf{e}_1, \dots, \mathbf{e}_n) : \mathbf{n} \mid \Gamma_1 \dots \Gamma_n, \kappa(\mathbf{n}, \mathbf{t}_1 \dots \mathbf{t}_n, \bar{\alpha})}$	$\bar{\alpha} \ \text{fresh}$
$(cast) \frac{\Pi \vdash \mathbf{e} : \mathbf{t} \mid \Gamma}{\Pi \vdash (\mathbf{n})\mathbf{e} : \mathbf{n} \mid \Gamma, \mathbf{n} \sim \mathbf{t}}$	

Fig. 8. Typing rules for classes