

A Core Calculus of Higher-Order Mixins and Classes^{*}

Lorenzo Bettini,¹ Viviana Bono,² and Silvia Likavec²

¹ Dipartimento di Sistemi ed Informatica, Università di Firenze, Via C. Lombroso 6/17,
50134 Firenze, Italy, bettini@dsi.unifi.it

² Dipartimento di Informatica, Università di Torino, C.so Svizzera 185,
10149 Torino, Italy, {bono,likavec}@di.unito.it

Abstract. This work presents an object-oriented calculus based on *higher-order* mixin construction via *mixin composition*, where some software engineering requirements are modeled in a formal setting allowing to prove the absence of *message-not-understood* run-time errors. Mixin composition is shown to be a valuable language feature enabling a cleaner object-oriented design and development. In what we believe being quite a general framework, we give directions for designing a programming language equipped with higher-order mixins, although our study is not based on any already existing object-oriented language.

1 Introduction

Recently, mixins are undergoing a renaissance (see, for example, [1, 5, 6]), due to their flexible nature of “incomplete” classes prone to be completed according to the programmer’s needs. Mixins [11, 15] are (sub)class definitions parameterized over a superclass and were introduced as an alternative to some forms of multiple inheritance [10, 17]. A mixin could be seen as a function that, given one class as an argument, produces another class, by adding or overriding certain sets of methods. The same mixin can be used to produce a variety of classes with the same functionality and behavior, since they all have the same sets of methods added and/or redefined. Also, the same mixin can sometimes be applied to the same class more than once, thus enabling incremental changes in the subclasses. The superclass definition is not needed at the time of writing the mixin definition. This minimizes the dependences between superclass and its subclasses, as well as between class implementors and end-users, thus improving modularity. The uniform extension and modification of classes is instead absent from the classical class-based languages. Moreover, in class-based languages, parentage is determined statically at compile time instead of at run-time.

In this work we extend the core calculus of classes and mixins of [8] with *higher-order* mixins. A mixin can: (*i*) be applied to a class to create a fully-fledged subclass; or (and this is the novelty with respect to [8]) (*ii*) be *composed* with another mixin to obtain yet another mixin with more functionalities. In Section 2.1 we present some uses of mixin inheritance and, in particular, we show that mixin composition enables a cleaner

^{*} This work has been partially supported by EU within the FET - Global Computing initiative, project MIKADO IST-2001-32222, DART project IST-2001-33477 and by MIUR project NAPOLI. The funding bodies are not responsible for any use that might be made of the results presented here.

modular object-oriented design. This paper presents a framework for the construction of composite mixins, and therefore of sophisticated class hierarchies, while keeping the good features of the original core calculus of [8]. In what we believe being quite a general framework, we give directions for designing a programming language equipped with higher-order mixins, although our study is not based on any already existing object-oriented language. The last section discusses how our approach attempts to cover and possibly ameliorate already existing calculi based on mixins, while keeping a good eye on the feasibility of the implementation of the calculus. The main difference with respect to the calculus of [8] is that we do not treat *protected* methods, being an orthogonal issue to higher-order mixins. Nevertheless, protected methods could be easily accounted for via (structural) subtyping as in the original calculus.

Our design decisions are strongly based on the choices that were made in [8]. Class hierarchies in a well-designed object-oriented program must not be fragile: if a superclass implementation changes but the specification remains intact, the implementors of the subclasses should not have to rewrite subclass implementations. This is only possible if object creation is modular. In particular, a subclass implementation should not be responsible for initializing inherited fields when a new object is created, since some of the inherited fields may be private and thus invisible to the subclass. Also, the definitions of inherited fields may change when the class hierarchy changes, making the subclass implementation invalid. Unlike many theoretical calculi for object-oriented languages, our calculus directly supports modular object construction. The mixin implementor only writes the local constructor for his own mixin. Mixin applications and compositions are reduced to generator functions that call all constructors in the inheritance chain in the correct order, producing a fully initialized object (see Section 3). Unlike some approaches to encapsulation in object calculi such as existential types, the levels of encapsulation describe *visibility*, and not merely *accessibility*. For example, even the names of private items are invisible outside the class in which they are defined. This seems to be a better approach since *no* information about data representation is revealed, not even the number and names of fields. One of the benefits of using visibility-based encapsulation is that no conflicts arise if both the superclass and the subclass declare a private field with the same name. Among other advantages, this allows the same mixin to be applied twice (see the example in Section 2.1). To ensure that mixin inheritance can be statically type checked, the calculus employs constrained parameterization. From each mixin definition the type system infers a constraint specifying to which classes the mixin can be applied so that the resulting subclass is type-safe. The constraint includes both *positive* (which methods the class must contain) and *negative* (which methods the class must not contain) information. New and redefined methods are distinguished in the mixin implementation. From the implementor's viewpoint, a new method may have arbitrary behavior, while the behavior of a redefined method must be "compatible" with that of the old method it replaces. Having this distinction in the syntax of our calculus helps mixin implementors avoid unintentional redefinitions of superclass methods and facilitates generation of the constraint for mixin's superclasses and for mixins that participate in mixin composition (see Section 4).

Expressions: $e ::= \text{const} \mid x \mid \lambda x.e \mid e_1 e_2 \mid \text{fix} \mid \text{ref} \mid ! \mid :=$
 $\mid \{x_i = e_i\}^{i \in I} \mid e.x \mid \mathbf{H} h.e \mid \text{new } e \mid e_1 \diamond e_2 \mid e_1 \bullet e_2$
 $\mid \text{classval}\langle v_g, \mathcal{M} \rangle \mid \text{mixinval}\langle v_m, \text{New}, \text{Redef}, \text{Expect} \rangle$
 $\mid \text{mixin}$
 $\quad \text{method } m_j = v_{m_j}; \quad (j \in \text{New})$
 $\quad \text{redefine } m_k = v_{m_k}; \quad (k \in \text{Redef})$
 $\quad \text{expect } m_i; \quad (i \in \text{Expect})$
 $\quad \text{constructor } v_c;$
 $\quad \text{end}$

Values: $v ::= \text{const} \mid x \mid \lambda x.e \mid \text{fix} \mid \text{ref} \mid ! \mid := \mid v \mid \{x_i = v_i\}^{i \in I}$
 $\mid \text{classval}\langle v_g, \mathcal{M} \rangle \mid \text{mixinval}\langle v_m, \text{New}, \text{Redef}, \text{Expect} \rangle$

Fig. 1. Syntax of the core calculus

2 Syntax of the Calculus

The starting point for our calculus is the core calculus of classes and mixins of Bono et al. [8] that, in turn, is based on *Reference ML* of Wright and Felleisen [19]. To this imperative calculus of records and functions, we add constructs for manipulating classes and mixins. The class and mixin related expressions are: `classval`, `mixin`, `mixinval`, `◇` (mixin application), `•` (mixin composition) and `new`. The novelties with respect to [8] are `mixinval` and `•` (mixin composition) to deal with *higher-order* mixins.

The majority of expressions (values) is standard and they are given in Figure 1. The only constructs that might need some explanation are the following:

- `ref`, `!`, `:=` are operators¹ for defining a reference to a value, for dereferencing a reference and for assigning a new value to a reference, respectively.
- $\{x_i = e_i\}^{i \in I}$ is a record and $e.x$ is the record selection operation (note that this corresponds to method selection in our calculus).
- h is a set of pairs $h := \{\langle x, v \rangle^*\}$ where x is a variable and v is a value (first components of the pairs are all distinct). We also have a concept of a *heap*, represented by h in the expression $\mathbf{H}h.e$, which is used for evaluating imperative side effects. In the expression $\mathbf{H}\langle x_1, v_1 \rangle \dots \langle x_n, v_n \rangle.e$, \mathbf{H} binds variables x_1, \dots, x_n in v_1, \dots, v_n and in e .
- `new` e uses generator v_g of the class value to which e evaluates to create a function that returns a new object, as described in Section 3.
- `classval` $\langle v_g, \mathcal{M} \rangle$ is a *class value*, and it is the result of mixin application. It is a pair, containing the function v_g , that is the generator for the class used to generate its instance objects, and the set \mathcal{M} of the indices of all the methods defined in the class. In our calculus method names are of the shape m_i , where i ranges over an index set, and are univocally identified by their index, i.e., $m_i = m_j$ if and only if $i = j$.

¹ Introducing `ref`, `!`, `:=` as operators rather than standard forms such as `ref e`, `!e`, `:=e1e2`, simplifies the definition of evaluation contexts and proofs of properties. As noted in [19], this is just a syntactic convenience, as is the curried version of `:=`.

- **mixin**
 method $m_j = v_{m_j};$ ($j \in New$)
 redefine $m_k = v_{m_k};$ ($k \in Redef$)
 expect $m_i;$ ($i \in Expect$)
 constructor $v_c;$
end

is a *mixin* expression, and it states the methods that are new, redefined, and expected in the mixin. More precisely, $m_j = v_{m_j}$ are definitions of the new methods, $m_k = v_{m_k}$ are method redefinitions that will replace the methods with the same name in the superclass, and m_i are method (names) that the superclass is expected to implement. Each method body v_{m_j} (respectively, v_{m_k}) is a function of the private *field* and of *self*, which will be bound to the newly created object at instantiation time. In method redefinitions, v_{m_k} is also a function of *next*, which will be bound to the corresponding old method from the superclass. The v_c value in the **constructor** clause is a function that returns a record of two components: the *fieldinit* value is used to initialize the private field; the *superinit* value is passed as an argument to the superclass constructor. When evaluating a mixin, v_c is used to build the generator as described in Section 3.

- $e_1 \diamond e_2$ denotes the application of mixin value e_1 to class value e_2 . Given the (super)class value e_2 as an “argument” to e_1 , it produces a new (sub)class value.
- $\text{mixinval}\langle v_m, New, Redef, Expect \rangle$ is a *mixin value*, and it is the result of a mixin evaluation. It is a tuple, containing one function and three sets of indices. The function v_m is the (partial) generator for the corresponding mixin. The sets *New*, *Redef*, and *Expect* contain the names of all methods defined in the mixin (new, redefined, and expected).
- $e_1 \bullet e_2$ is a composition of two mixin values e_1 and e_2 . It produces a new mixin value taking components from both e_1 and e_2 . The resulting mixin can be applied to class values to produce new classes, as well as composed with other mixin values to produce new composite mixins.

As in [8], we define the root of the class hierarchy, class *Object*, as a predefined class value: $Object \triangleq \text{classval}\langle \lambda_.\lambda_.\{\}, [] \rangle$. The root class is necessary so that all other classes can be treated uniformly and it is the only class value that is not obtained as a result of mixin application. The calculus can then be simplified by assuming that any user-defined class that does not need a superclass is obtained by applying a mixin containing all of the class method definitions to *Object*. For the sake of clarity, in the following examples we will avoid the explicit mixin application to *Object*.

2.1 An example of mixin inheritance

In this section, we present a simple example that shows how mixins can be implemented and used in our calculus and explain some of the uses of mixin application and mixin composition. For readability, the example uses functions with multiple arguments even though they are not formalized explicitly in the calculus.

In the following, we give the definitions of **Encrypted** mixin and **Compress** mixin that implement encryption and compression functionality on top of any stream class, respectively. Note that the class to which the mixin is applied may have more methods

than expected by the mixin. For example, `Encrypted` can be applied to `Socket` \diamond `Object`, even though `Socket` \diamond `Object` has other methods besides `read` and `write`. The mixin `Random` allows random access to any stream class, thus we can build a random access file class with the mixin application `Random` \diamond `FileStream`.

```

let FileStream = mixin
  method write = ...
  method read = ...
end in

let Socket = mixin
  method write = ...
  method read = ...
  method IPaddress = ...
end in

let Random = mixin
  method lseek = ...
  expect write;
  expect read;
end in

let Encrypted =
  mixin
  redefine write =  $\lambda$  key.  $\lambda$  self.  $\lambda$  next.  $\lambda$  data. next (encrypt(data, key));
  redefine read =  $\lambda$  key.  $\lambda$  self.  $\lambda$  next.  $\lambda$  _ . decrypt(next (), key);
  constructor  $\lambda$  (key, arg). {fieldinit=key, superinit=arg};
end in

let Compress =
  mixin
  redefine write =  $\lambda$  level.  $\lambda$  self.  $\lambda$  next.  $\lambda$  data. next (compress(data, level));
  redefine read =  $\lambda$  level.  $\lambda$  self.  $\lambda$  next.  $\lambda$  _ . uncompress(next (), level);
  constructor  $\lambda$  (level, arg). {fieldinit=level, superinit=arg};
end in ...

```

From the definition of `Encrypted`, the type system infers the types of the methods that the mixin wants to redefine. These are the constraints that must be satisfied by any class to which `Encrypted` is applied. The class must contain `write` and `read` methods whose types must be supertypes of those given to `write` and `read`, respectively, in the definition of `Encrypted`. In `RandomFile` such methods are declared as *expected* and they are used within the method `lseek`. Once again the type system infers their types according to how they are used in `lseek`.

To create an encrypted stream class, one must apply the `Encrypted` mixin to an existing stream class. For example, `Encrypted` \diamond `FileStream` is an encrypted file class. The power of mixins can be seen when we apply `Encrypted` to a family of different streams. For example, we can construct `Encrypted` \diamond `Socket`, which is a class that encrypts data communicated over a network. In addition to single inheritance, we can express many uses of multiple inheritance by applying more than one mixin to a class. For example, `PGPSign` \diamond `UUEncode` \diamond `Encrypted` \diamond `Compress` \diamond `FileStream` produces a class of files that are compressed, then encrypted, then uuencoded, then signed. In addition, mixins can be used for forms of inheritance that are not possible in most single and multiple inheritance-based systems. In the above example, the result of applying `Encrypted` to a stream satisfies the constraint required by `Encrypted` itself, therefore, we can apply `Encrypted` more than once: `Encrypted` \diamond `Encrypted` \diamond `FileStream` is a class of files that are encrypted twice. In our system, class private fields do not conflict even if they have the same name, so each application of `Encrypted` can have its own encryption key.

Mixin composition further enhances the (re)usability of classes and mixins and enables better modular programming design, by exploiting software composition at a higher level. For example, the programmer is able to build a customized library of reusable mix-

ins starting from existing mixins: one can create the new mixin `2Encrypt = Encrypted • Encrypted`, instead of always applying the mixin `Encrypted` twice to every stream class in her program. This also enables consistency: if in the future the definition of the mixin `2Encrypt` has to be extended, e.g., by also exploiting UU encoding, then by changing only the definition of `2Encrypt`, with an additional mixin composition, it is guaranteed that all the functions that used `2Encrypt` will use the new version. Moreover, construction of mixins can be delegated to different parts of the program (thus exploiting modular programming) and the resulting mixins can then be assembled in order to build a class. For instance, the following code delegates the construction of mixins for encryption and compression to two functions, and then assembles the returned mixins for later use:

```
let m1 = build_compression() in let m2 = build_encryption() in
let m = m1 • m2 in (new(m ◊ FileStream)).write("foo")
```

The function `build_compression` returns a specific mixin according to user's requests: it can return a simple `Compress` mixin, or a more elaborate `UUEncode • Compress` mixin. Similarly, `build_encryption`, instead of simply returning a mixin `Encrypted`, returns the composition `PGPSign • Encrypted`. All these enhanced modular composition functionalities, supported by mixin composition, would not be directly provided by simple mixin application.

3 Operational Semantics

The operational semantics extends the one of the core calculus of classes and mixins, [8], and therefore exploits the *Reference ML* of Wright and Felleisen treatment of side-effects [19]. We give the reduction rules in Figures 2 and 4. To abstract from a precise set of constants, we only assume the existence of a partial function $\delta : Const \times ClosedVal \rightarrow ClosedVal$ that interprets the application of functional constants to closed values and yields closed values. In Figure 2, R are the *reduction contexts* [13, 14, 18]. Reduction contexts are necessary to provide a minimal relative linear order among the creation, dereferencing and updating of heap locations, since side effects need to be evaluated in a deterministic order. Their definition can be found in Figure 3. We assume the reader is familiar with the treatment of imperative side-effects via reduction contexts and we refer to [19, 8] for a description of the related rules.

(*new*) rule is responsible for instantiating new objects from class definitions. The resulting function can be thought of as the composition of two functions: $fix \circ g$. First, the generator g is applied to an argument v , thus creating a function from *self* to a record of methods. Afterwards, the fixed-point operator fix (following [12]) is applied to bind *self* in method bodies and create a recursive record. The resulting record is a fully formed object that could be returned to the user.

(*mixval*) rule turns a mixin expression into a mixin *value*. A mixin value consists of a mixin generator Gen_m and of the sets of mixin method names (new, redefined, and expected; we recall that names are identified with their indices, as said in Section 2). Gen_m is a sort of compiled (equivalent) version of the mixin expression: given the parameter for the mixin constructor c , it returns a record containing a (partial) object generator `gen`, and the argument `superinit` for the (future) superclass constructor. We recall that c is a function of one argument which returns a record of two components: one

$$\begin{aligned}
& \text{mixinval}\langle g_1, New_1, Redef_1, Expect_1 \rangle \bullet \text{mixinval}\langle g_2, New_2, Redef_2, Expect_2 \rangle \rightarrow \\
& \text{mixinval}\langle Gen, New_1 \cup New_2, (Redef_1 \cup Redef_2) - New_2, \\
& \quad (Expect_1 - (New_2 \cup Redef_2)) \cup (Expect_2 - Redef_1) \rangle \\
Gen & \triangleq \lambda x. \\
& \text{let } leftrec = g_1(x) \text{ in} \\
& \text{let } rightrec = g_2(leftrec.superinit) \text{ in} \\
& \text{let } leftgen = leftrec.gen \text{ in} \\
& \text{let } rightgen = rightrec.gen \text{ in} \\
& \left(\begin{array}{l}
gen = \lambda self. \\
\left(\begin{array}{l}
m_{j_1} = \lambda y. (leftgen \ self).m_{j_1} \ y \quad j_1 \in New_1 \\
m_{j_2} = \lambda y. (rightgen \ self).m_{j_2} \ y \quad j_2 \in New_2 - Redef_1 \\
m_{j_3} = \lambda y. (leftgen \ self).m_{j_3} \ (rightgen \ self).m_{j_3} \ y \quad j_3 \in Redef_1 \cap New_2 \\
m_{k_1} = \lambda y. (leftgen \ self).m_{k_1} \ y \quad k_1 \in Redef_1 - (New_2 \cup Redef_2) \\
m_{k_2} = \lambda next. (leftgen \ self).m_{k_2} \ ((rightgen \ self).m_{k_2} \ next) \quad k_2 \in Redef_1 \cap Redef_2 \\
m_{k_3} = \lambda y. (rightgen \ self).m_{k_3} \ y \quad k_3 \in Redef_2 - Redef_1
\end{array} \right) \\
superinit = rightrec.superinit
\end{array} \right)
\end{aligned}$$

Fig. 4. Reduction rule (mixcomp) for mixin composition

Here mixin values are made explicit to deal smoothly with mixin composition. For all the methods, the method bodies are wrapped inside $\lambda y. \dots y$ to delay evaluation in our call-by-value calculus.

(*mixapp*) rule evaluates the application of a mixin value to a class value, performing mixin-based inheritance. A mixin value $\text{mixinval}\langle Gen_m, New, Redef, Expect \rangle$ is applied to a (super)class value $\text{classval}\langle g, \mathcal{M} \rangle$, where g is the object generator of the class and \mathcal{M} is the set of all method names defined in the superclass. The resulting class value is $\text{classval}\langle Gen, New \cup \mathcal{M} \rangle$, where Gen is the generator function for the subclass, and $New \cup \mathcal{M}$ lists all its method names. Using a class generator delays full inheritance resolution until object instantiation time when *self* becomes available. The class generator takes a single argument x , which is used by the mixin generator, and returns a function from *self* to a record of methods. When the fixed-point operator is applied to the function returned by the class generator, it produces a recursive record of methods representing a new object (see rule (*new*)). Gen first calls $Gen_m(x)$ to compute the mixin object generator *mixingen*, a function from *self* to a record of mixin methods, and the parameter *mixingen.superinit* to be passed to the superclass generator g , that, in turn, returns a function *supergen* from *self* to a record of superclass methods. Gen results to be a function of *self* that returns a record containing *all* the methods — from both the mixin and the superclass. All methods of the superclass that are not redefined by the mixin, m_i where $i \in \mathcal{M} - Redef$, are *inherited* by the subclass: they are taken intact from the superclass’s “object” (*supergen self*). These methods m_i include all the methods that are expected by the mixin (as checked by the type system). Methods m_j defined by the mixin are taken intact from the mixin’s “object” (*mixingen self*). As for *redefined* methods m_k , *next* is bound to (*supergen self*). m_k in Gen . Notice that at this stage, all methods have already received a binding for the private field. The variable *self* is passed all along in all method forms, in such a way that the host object is bound appropriately everywhere at object creation time.

(*mixcomp*) (Fig. 4) composes two mixins in order to produce a new mixin. The two mixins may partially complete each others' definitions, providing (some of) the missing components. Let us denote the mixin composition by $e_1 \bullet e_2$ and the resulting mixin by e . When composing two mixins, it is necessary to determine which sets of new/redefined/expected methods the new mixin e will have. Our design decision is as follows: the mixin e_2 acts as a “superclass” for e_1 (mirroring mixin application order), and, in particular, some of e_1 methods may override some of e_2 methods. Therefore, all the new methods of the mixin e_1 (New_1) are inserted in the resulting mixin e , while only the new methods of e_2 that are not redefined by e_1 ($j_2 \in New_2 - Redef_1$) become part of the new mixin. Notice that the type rule for mixin composition (*mixin comp*) (Figure 6) must check that no name clashes between new methods of e_1 and any method of e_2 take place. This decision is in line with a good object-oriented design principle of not confusing method redefinitions and name clashes. Therefore, an error is signaled at compile time and not at runtime. As far as redefined methods are concerned, the situation is more complex: the methods specified as redefined in e_1 can override some new methods from e_2 , some redefined methods from e_2 , and (even if only virtually) some of the expected methods from e_2 .

- If a method m_{j_3} in e_1 redefines a method defined in e_2 ($j_3 \in Redef_1 \cap New_2$), then the overriding is completed and m_{j_3} becomes a new method in the resulting mixin e , after binding its *next* to e_2 's implementation of m_{j_3} ;
- If e_1 redefines a method m_{k_2} that, in turn, is redefined by e_2 ($k_2 \in Redef_1 \cap Redef_2$), then this method is still a redefined method in e . Since e_1 “overrides” e_2 , and m_{k_2} 's implementation of e_1 redefines that of e_2 , the *next* in the implementation of e_1 is bound to the implementation of e_2 , and the *next* in the implementation of e_2 is not bound, since it will be bound during future mixin composition or mixin application. This means that the redefinition of a method m_{k_2} by means of e_2 is delayed (while e_1 has already performed its “internal” redefinition of m_{k_2} over e_2);
- If e_1 redefines a method that is expected in e_2 , then this method will become a redefined method in e , so it will not appear among the expected methods of e , but it will be a method that e is willing to redefine.

Apart from the above examined methods, method redefinitions that are still present as method redefinitions in the resulting mixin e are: (i) the ones from e_2 that are not redefined by e_1 ($k_3 \in Redef_2 - Redef_1$); (ii) the ones from e_1 that are not defined in e_2 and hence not “overriding” anything yet ($k_1 \in Redef_1 - (New_2 \cup Redef_2)$).

Finally, new and redefined methods from e_2 can provide some of the definitions that the mixin e_1 expects; in that case, such methods expected by e_1 do not appear anymore in the expected method set of e .

The generator of the new mixin is a combination of the generators of e_1 and e_2 : since e_1 is considered to be the “subclass”, the parameter x is passed to g_1 , and g_2 receives as a parameter the *superinit* returned by $g_1(x)$; the *superinit* field of the record returned by the generator of the new mixin is set to $g_2(g_1(x).superinit).superinit$. This strategy for building the new mixin generator corresponds to serializing the call of the two constructors similarly to what happens in standard object-oriented languages. Notice that this is consistent with the type $\text{mixin}\langle \gamma_{b_2}, \gamma_{d_1}, \Sigma_{new}, \Sigma_{red}, \Sigma_{exp}, \Sigma_{old} \rangle$ assigned to the new mixin by the type rule (*mixin comp*) (Figure 6).

4 Type System

In addition to functional, record, and reference types of *Reference ML* type system, our type system has class-types and mixin-types.

The types in our system are the following:

$$\tau ::= \iota \mid \tau_1 \rightarrow \tau_2 \mid \tau \text{ ref} \mid \{m_i : \tau_{m_i}\}^{i \in I} \mid \text{class}\langle \tau, \Sigma_b \rangle \mid \text{mixin}\langle \tau_1, \tau_2, \Sigma_{new}, \Sigma_{red}, \Sigma_{exp}, \Sigma_{old} \rangle$$

where ι is a constant type, \rightarrow is the functional type operator, $\tau \text{ ref}$ is the type of locations containing a value of type τ . The other type forms are described below.

Σ (possibly with a subscript) denotes a record type of the form $\{m_i : \tau_{m_i}\}^{i \in I}$. The set of indexes I (where $I \subseteq \mathbb{N}$) is often omitted when it is not relevant. A record type can be viewed as a set of pairs *label:type* where labels are pairwise disjoint (Σ_1 and Σ_2 are considered *equal*, denoted by $\Sigma_1 = \Sigma_2$, if they differ only in the order of their elements).

Thus notations and operations on sets are easily extended to record types as in the following definitions:

- if $m_i : \tau_{m_i} \in \Sigma$ we say that the *subject* m_i *occurs* in Σ (with type τ_{m_i}). $\text{Subj}(\Sigma)$ denotes the set of all subjects occurring in Σ ;
- $\Sigma_1 \cup \Sigma_2$ is the standard set union (used only on Σ_1 and Σ_2 such that $\text{Subj}(\Sigma_1) \cap \text{Subj}(\Sigma_2) = \emptyset$, in order to guarantee that $\Sigma_1 \cup \Sigma_2$ is a record type);
- $\Sigma_1 - \Sigma_2$ is the standard set difference;
- $\Sigma_1 / \Sigma_2 = \{m_i : \tau_{m_i} \mid m_i : \tau_{m_i} \in \Sigma_1 \wedge m_i \text{ occurs in } \Sigma_2\}$.

The definitions of typing environments Γ and of typing judgments are standard. Our type system supports *structural subtyping* ($<:$ relation) along with a subsumption rule (*sub*). The subtyping rules are shown in Appendix A. Since subtyping on references is unsound and we wish to keep subtyping and inheritance completely separate, we have only the basic subtyping rules for function and record types. Subtyping only exists at the object level, and is not supported for class or mixin types.

In the class type $\text{class}\langle \gamma, \Sigma_b \rangle$, γ is the type of the generator's argument and $\Sigma_b = \{m_i : \tau_{m_i}\}$ is a record type representing *self*.

In the mixin type $\text{mixin}\langle \gamma_b, \gamma_d, \Sigma_{new}, \Sigma_{red}, \Sigma_{exp}, \Sigma_{old} \rangle$

- γ_b is the expected argument type of the superclass generator,
- γ_d is the exact argument type of the mixin generator,
- $\Sigma_{new} = \{m_j : \tau_{m_j}^\dagger\}$ are the exact types of the new methods introduced by the mixin,
- $\Sigma_{red} = \{m_k : \tau_{m_k}^\dagger\}$ are the exact types of the methods redefined by the mixin,
- $\Sigma_{exp} = \{m_i : \tau_{m_i}^\dagger\}$ are the types of the methods that are neither defined nor redefined by the mixin but expected to be supported by a superclass to which the mixin will be applied to, or by a mixin to which it will be composed with,
- $\Sigma_{old} = \{m_k : \tau_{m_k}^\dagger\}$ are the types assumed for the old bodies of the methods redefined by the mixin.

We report in Figure 5 the majority of typing rules regarding classes and mixins that are syntactic variations of those presented in [8] (we refer the reader to that paper for a comment about these rules; other type rules are given in Appendix A). We only comment upon the rules related to mixin forms. The rules (*mixin*) and (*mixin val*)

assign the same type to their respective expressions, although deduced in a different way. In the rule (*mixin app*), Σ_b contains the type signatures of all methods supported by the superclass to which the mixin is applied and Σ_b/Σ_{red} are the superclass methods redefined by the mixin (the superclass may have more methods than those required by the mixin constraints). The premises of the rule (*mixin app*) are the following:

- i) $\Sigma_b <: (\Sigma_{exp} \cup \Sigma_{old})$ requires the actual types of the superclass methods to be subtypes of those expected by the mixin.
- ii) $\Sigma_{red} <: \Sigma_b/\Sigma_{red}$ requires that the types of the actual implementations of methods in the superclass (which may belong to a subtype of the Σ_{old} , from the above constraint) are supertypes of the ones redefined in the mixin. Thus, the types of the methods redefined by the mixin (Σ_{red}) will be subtypes of the superclass methods with the same name.
- iii) $Subj(\Sigma_b) \cap Subj(\Sigma_{new}) = \emptyset$ guarantees that no name clash will take place during the mixin application.

Intuitively, the above constraints insure that all the actual method bodies of the newly created subclass are at least as “good” as expected. The resulting class, of type $\text{class}\langle\gamma_d, \Sigma_d\rangle$, contains the signatures of all the methods forming the new class, created as the result of mixin application. Σ_{red} and Σ_{new} are methods defined by the mixin, whereas $\Sigma_b - (\Sigma_b/\Sigma_{red})$ are the methods inherited directly from the superclass. Let us observe that, for any well typed mixin, $Subj(\Sigma_{red}) = Subj(\Sigma_{old})$, therefore for any record type Σ , $\Sigma/\Sigma_{red} = \Sigma/\Sigma_{old}$.

Now we concentrate on the rule for mixin composition (*mixin comp*) given in Figure 6, which is the main topic of the paper. Since e_2 acts as the “superclass” of e_1 , e_1 will pass the argument of type γ_{b_1} to the constructor of the superclass e_2 , that expects an argument of type γ_{d_2} for its constructor. Therefore, we require that $\gamma_{b_1} <: \gamma_{d_2}$ (condition (c_1)).

The mixin e_1 is allowed to redefine methods defined by e_2 , expected by e_2 , and methods that are in turn redefined by e_2 . In all the cases we must check that the redefinition (and the expectation about the old method in the superclass) is type safe (conditions (c_2) , (c_3) and (c_4)). If e_1 redefines a method m_k that is in turn redefined by e_2 , then we will put the redefined type of m_k from e_1 in Σ_{red} and the old one from e_2 in Σ_{old} . This is consistent with the view that the new mixin will contain m_k with the body from e_1 (with its *next* bound to e_2 ’s implementation, while in m_k ’s body from e_2 *next* remains still unbound, as the method m_k can be further redefined, see Section 3). If e_1 redefines, instead, an expected method of e_2 , then that method will not appear in Σ_{exp} , but the redefined type and the old type, as inferred from e_1 , will appear in Σ_{red} and Σ_{old} , respectively.

Conditions (c_5) and (c_6) check whether e_2 can provide methods (either defined or redefined) that are expected by e_1 . If such a method is provided, then it will not appear in Σ_{exp} . In case both e_1 and e_2 expect the same method, the types with which such method is expected must be comparable (condition (c_7)); the method will then appear in Σ_{exp} with the smaller type.

Finally, condition (c_8) checks that no name clash occurs among methods defined by e_1 and those defined/redefined/expected by e_2 . This decision is in line with a good object-oriented design principle of not confusing method redefinitions and name clashes.

$$\begin{array}{c}
\frac{\Gamma \vdash g : \gamma \rightarrow \{m_i : \tau_{m_i}\}^{i \in \mathcal{M}} \rightarrow \{m_i : \tau_{m_i}\}^{i \in \mathcal{M}}}{\Gamma \vdash \text{classval}\langle g, \mathcal{M} \rangle : \text{class}\langle \gamma, \{m_i : \tau_{m_i}\}^{i \in \mathcal{M}} \rangle} \text{ (class val)} \quad \frac{\Gamma \vdash e : \text{class}\langle \gamma, \{m_i : \tau_{m_i}\} \rangle}{\Gamma \vdash \text{new } e : \gamma \rightarrow \{m_i : \tau_{m_i}\}} \text{ (instantiate)} \\
\\
\begin{array}{l}
\text{(New) For } j \in \text{New}: \Gamma \vdash v_{m_j} : \eta \rightarrow \Sigma \rightarrow \tau_{m_j}^\dagger \\
\text{(Redef) For } k \in \text{Redef}: \Gamma \vdash v_{m_k} : \eta \rightarrow \Sigma \rightarrow \tau_{m_k}^\dagger \rightarrow \tau_{m_k}^\dagger \\
\text{(Constr)} \quad \Gamma \vdash c : \gamma_d \rightarrow \{\text{fieldinit} : \eta, \text{superinit} : \gamma_b\}
\end{array} \\
\hline
\Gamma \vdash \left(\begin{array}{l} \text{mixin} \\ \text{method } m_j = v_{m_j}; \\ \text{redefine } m_k = v_{m_k}; \\ \text{expect } m_i; \\ \text{constructor } c; \\ \text{end} \end{array} \right) : \text{mixin}\langle \gamma_b, \gamma_d, \Sigma_{\text{new}}, \Sigma_{\text{red}}, \Sigma_{\text{exp}}, \Sigma_{\text{old}} \rangle \quad \begin{array}{l} j \in \text{New} \\ k \in \text{Redef} \\ i \in \text{Expect} \end{array} \text{ (mixin)} \\
\\
\frac{\Gamma \vdash g : \gamma_d \rightarrow \{\text{gen} : \Sigma \rightarrow \{m_j : \tau_{m_j}^\dagger, m_k : \tau_{m_k}^\dagger \rightarrow \tau_{m_k}^\dagger\}^{j \in \text{New}, k \in \text{Redef}}, \text{superinit} : \gamma_b\}}{\Gamma \vdash \text{mixINVAL}\langle g, \text{New}, \text{Redef}, \text{Expect} \rangle : \text{mixin}\langle \gamma_b, \gamma_d, \Sigma_{\text{new}}, \Sigma_{\text{red}}, \Sigma_{\text{exp}}, \Sigma_{\text{old}} \rangle} \text{ (mixin val)} \\
\\
\begin{array}{l}
\Sigma = \Sigma_{\text{new}} \cup \Sigma_{\text{red}} \cup \Sigma_{\text{exp}} \\
\text{where } \Sigma_{\text{new}} = \{m_j : \tau_{m_j}^\dagger\}, \Sigma_{\text{red}} = \{m_k : \tau_{m_k}^\dagger\}, \Sigma_{\text{exp}} = \{m_i : \tau_{m_i}^\dagger\}, \Sigma_{\text{old}} = \{m_k : \tau_{m_k}^\dagger\} \\
\tau_{m_i}^\dagger \text{ and } \tau_{m_k}^\dagger \text{ are inferred from method bodies and } i \in \text{Expect}
\end{array} \\
\\
\begin{array}{l}
\Gamma \vdash e_1 : \text{mixin}\langle \gamma_b, \gamma_d, \Sigma_{\text{new}}, \Sigma_{\text{red}}, \Sigma_{\text{exp}}, \Sigma_{\text{old}} \rangle \\
\Gamma \vdash e_2 : \text{class}\langle \gamma_c, \Sigma_b \rangle \\
\Gamma \vdash \gamma_b <: \gamma_c \\
\Gamma \vdash \Sigma_b <: (\Sigma_{\text{exp}} \cup \Sigma_{\text{old}}) \\
\Gamma \vdash \Sigma_{\text{red}} <: \Sigma_b / \Sigma_{\text{red}} \\
\text{Subj}(\Sigma_b) \cap \text{Subj}(\Sigma_{\text{new}}) = \emptyset
\end{array} \\
\hline
\Gamma \vdash e_1 \diamond e_2 : \text{class}\langle \gamma_d, \Sigma_d \rangle \quad \text{(mixin app)} \\
\\
\text{where } \Sigma_d = \Sigma_{\text{new}} \cup \Sigma_{\text{red}} \cup (\Sigma_b - \Sigma_b / \Sigma_{\text{red}})
\end{array}$$

Fig. 5. Typing rules for class and mixin-related forms

Our system is proved sound, in the sense that “every well-typed program cannot go wrong”, which implies the absence of *message-not-understood* runtime errors.

Lemma 1 (Subject Reduction). *If $\Gamma \vdash e : \tau$ and e evaluates to e' , then $\Gamma \vdash e' : \tau$.*

We consider *programs*, which are closed terms. Then, we introduce *faulty programs* which are a way to approximate the concept of reaching a “stuck state” during the evaluation process, and prove that if the evaluation for a given program p does not diverge, then either it returns a value, or p reduces to a faulty program. By using the subject reduction property and proving that faulty programs are not typable, we show that if a program has a type in our system, then it evaluates to a value, under the condition that the program does not diverge.

Theorem 1 (Soundness). *Let p be a program: if $\varepsilon \vdash p : \tau$ then either the evaluation for p diverges, or p evaluates to a value v and $\varepsilon \vdash v : \tau$ (ε stands for the empty typing environment).*

$$\begin{array}{l}
\Gamma \vdash e_1 : \text{mixin}(\gamma_{b_1}, \gamma_{d_1}, \Sigma_{new}^1, \Sigma_{red}^1, \Sigma_{exp}^1, \Sigma_{old}^1) \\
\Gamma \vdash e_2 : \text{mixin}(\gamma_{b_2}, \gamma_{d_2}, \Sigma_{new}^2, \Sigma_{red}^2, \Sigma_{exp}^2, \Sigma_{old}^2) \\
(c_1) \Gamma \vdash \gamma_{b_1} <: \gamma_{d_2} \\
(c_2) \Gamma \vdash \tau_{m_{k_1}}^\downarrow <: \tau_{m_{j_2}}^{\prime\downarrow} <: \tau_{m_{k_1}}^\uparrow \quad \text{if } k_1 = j_2 \\
(c_3) \Gamma \vdash \tau_{m_{k_1}}^\downarrow <: \tau_{m_{k_2}}^{\prime\downarrow} <: \tau_{m_{k_1}}^\uparrow \quad \text{if } k_1 = k_2 \\
(c_4) \Gamma \vdash \tau_{m_{k_1}}^\downarrow <: \tau_{m_{i_2}}^{\prime\downarrow} <: \tau_{m_{k_1}}^\uparrow \quad \text{if } k_1 = i_2 \\
(c_5) \Gamma \vdash \tau_{m_{j_2}}^{\prime\downarrow} <: \tau_{m_{i_1}}^\uparrow \quad \text{if } i_1 = j_2 \\
(c_6) \Gamma \vdash \tau_{m_{k_2}}^{\prime\downarrow} <: \tau_{m_{i_1}}^\uparrow \quad \text{if } i_1 = k_2 \\
(c_7) \Gamma \vdash \tau_{m_{i_2}}^{\prime\downarrow} <: \tau_{m_{i_1}}^\uparrow \vee \Gamma \vdash \tau_{m_{i_1}}^\uparrow <: \tau_{m_{i_2}}^{\prime\downarrow} \quad \text{if } i_1 = i_2 \\
(c_8) \text{Subj}(\Sigma_{new}^1) \cap (\text{Subj}(\Sigma_{new}^2) \cup \text{Subj}(\Sigma_{red}^2) \cup \text{Subj}(\Sigma_{exp}^2)) = \emptyset \\
\hline
\Gamma \vdash e_1 \bullet e_2 : \text{mixin}(\gamma_{b_2}, \gamma_{d_1}, \Sigma_{new}, \Sigma_{red}, \Sigma_{exp}, \Sigma_{old}) \quad (\text{mixin comp})
\end{array}$$

$$\text{where } \begin{array}{l}
\Sigma_{new}^1 = \{m_{j_1} : \tau_{m_{j_1}}^\downarrow\}, \Sigma_{new}^2 = \{m_{j_2} : \tau_{m_{j_2}}^{\prime\downarrow}\}, \Sigma_{red}^1 = \{m_{k_1} : \tau_{m_{k_1}}^\downarrow\}, \Sigma_{red}^2 = \{m_{k_2} : \tau_{m_{k_2}}^{\prime\downarrow}\} \\
\Sigma_{exp}^1 = \{m_{i_1} : \tau_{m_{i_1}}^\uparrow\}, \Sigma_{exp}^2 = \{m_{i_2} : \tau_{m_{i_2}}^{\prime\uparrow}\}, \Sigma_{old}^1 = \{m_{k_1} : \tau_{m_{k_1}}^\uparrow\}, \Sigma_{old}^2 = \{m_{k_2} : \tau_{m_{k_2}}^{\prime\uparrow}\}
\end{array}$$

$$\begin{array}{l}
\Sigma_{new} = \Sigma_{new}^1 \cup (\Sigma_{new}^2 - \Sigma_{new}^2 / \Sigma_{red}^1) \cup \Sigma_{red}^1 / \Sigma_{new}^2 \\
\Sigma_{red} = (\Sigma_{red}^1 - \Sigma_{red}^1 / \Sigma_{new}^2) \cup (\Sigma_{red}^2 - \Sigma_{red}^2 / \Sigma_{red}^1) \\
\Sigma_{old} = (\Sigma_{old}^1 - (\Sigma_{old}^1 / \Sigma_{new}^2 \cup \Sigma_{old}^1 / \Sigma_{old}^2)) \cup \Sigma_{old}^2 \\
\Sigma_{exp} = (\Sigma_{exp}^1 - (\Sigma_{exp}^1 / \Sigma_{new}^2 \cup \Sigma_{exp}^1 / \Sigma_{red}^2 \cup \Sigma_{exp}^1 / \Sigma_{exp}^2)) \cup (\Sigma_{exp}^2 - (\Sigma_{exp}^2 / \Sigma_{red}^1 \cup \Sigma_{exp}^2 / \Sigma_{exp}^1)) \cup \Sigma_{min} \\
\Sigma_{min} = \{m_i : \min\{\tau_{m_i}^\uparrow, \tau_{m_i}^{\prime\uparrow}\} \mid m_i : \tau_{m_i}^\uparrow \in \Sigma_{exp}^1, m_i : \tau_{m_i}^{\prime\uparrow} \in \Sigma_{exp}^2\}
\end{array}$$

Fig. 6. Typing rule for mixin composition

The metatheory for the present system, and in particular the subject reduction property, are extensions of the ones in [7] (Chapter 9). The formal definitions and properties were analyzed in detail, but they are omitted here for lack of space.

5 Related and Future Work

Bracha and Cook extend Modula-3 with mixins in [11] (one of the seminal papers on mixins). The novelty is seeing object types as mixins, which either explicitly state the modifications to the superclass, or are obtained as a result of mixin composition. The left-hand mixin has a “priority” and the composition is not explicitly written in order to ensure upward compatibility with the existing language. Instead, we think that making the composition explicit (as it is in our calculus) makes the programmer aware of how software components are composed, thus providing more control over the behavior of the program.

Flatt et al. [16] extend a subset of sequential Java called CLASSICJAVA with mixins and call it MIXEDJAVA. Mixins use their *inheritance interface* to specify how the inherited methods are extended and/or overridden. Existing mixins can be combined in order to produce new composite mixins. As in our calculus, the left-hand mixin has the “precedence” over the right-hand mixin. Composition is well-defined only if the right-hand mixin implements the left-hand mixin inheritance interface (i.e., the right-hand mixin is required to provide all the methods expected by the left-hand one). In this respect, our approach is more oriented to code composition, in that the new composite mixin is still allowed to have expected methods not yet resolved. The duplication of method names is resolved at run-time with the run-time context information provided

by the current *view* of the object (represented as a chain of mixins). We preferred to deal with name clashes statically by forcing the programmer to solve them, since automatic handling of such ambiguities may lead to unexpected behavior at run-time.

Boudol [9] extends *Reference ML* [19] with records and *let rec* operator. This enriched ML leads to a theoretically solid treatment of mixins, which are seen as class transformers. The main difference between the two calculi seems to be in the way references to fields are created. In our calculus these are created at class creation time, when mixin application is evaluated, whereas in the calculus of Boudol they are created at class instantiation time, i.e., when an object is created.

Ancona and Zucca [2, 3] give a formal model for mixin modules. A mixin is a function from input to output components, and they characterize axiomatically the operators for composing mixins in order to obtain higher-order mixins. In [1] they present JAM, an extension of Java supporting mixins, but not mixin composition.

The first future direction we would like to take is to study the extension of this calculus where not only classes can be instantiated but also mixins. A mixin can be seen as an “incomplete” class, so its objects are *incomplete* objects, which can be: (i) used, with respect to the methods that are already completed; (ii) enriched, using a method addition operation. This is an integration between a class-based and an object-based language, yielding a language suited for rapid prototyping (a first version of the calculus of incomplete objects can be found in [4]). Finally, higher-order mixins seem to be a natural feature to be added to MoMi [5], a coordination language where object-oriented mobile code is exchanged among the nodes of a network.

References

1. D. Ancona, G. Lagorio, and E. Zucca. Jam - a smooth extension of java with mixins. In *Proc. ECOOP 2000 - Object-Oriented Programming, 14th European Conference*, pages 154–178. LNCS 1850, Springer-Verlag, 2000.
2. D. Ancona and E. Zucca. An algebraic approach to mixins and modularity. In *Proc. Algebraic and Logic Programming (ALP)*, pages 179–193. LNCS 1139, Springer-Verlag, 1996.
3. D. Ancona and E. Zucca. An algebra of mixin modules. In *Proc. Recent Trends in Algebraic Development Techniques, 12th Int. Workshop, WADT'97*, number 1376 in LNCS, pages 92–106. Springer, 1997.
4. L. Bettini, V. Bono, and S. Likavec. A Core Calculus of Mixin-Based Incomplete Objects. Tech. report. Submitted for publication.
5. L. Bettini, V. Bono, and B. Venneri. Coordinating Mobile Object-Oriented Code. In F. Arbarb and C. Talcott, editors, *Proc. of Coordination Models and Languages*, number 2315 in LNCS, pages 56–71. Springer, 2002.
6. L. Bettini, V. Bono, and B. Venneri. Subtyping Mobile Classes and Mixins. In *Proc. of Int. Workshops on Foundations of Object-Oriented Languages, FOOL 10*, 2003.
7. V. Bono. *Type Systems for the Object Oriented Paradigm*. PhD thesis, Università di Torino, 1999.
8. V. Bono, A. Patel, and V. Shmatikov. A core calculus of classes and mixins. In *Proc. ECOOP '99*, pages 43–66. LNCS 1628, Springer-Verlag, 1999.
9. G. Boudol. The recursive record semantics of objects revised. In *Proc. ESOP '01*, pages 269–283. LNCS 2028, Springer-Verlag, 2001.
10. N. Boyen, C. Lucas, and P. Steyaert. Generalized mixin-based inheritance to support multiple inheritance. Technical Report vub-prog-tr-94-12, Vrije Universiteit Brussel, 1994.

11. G. Bracha and W. Cook. Mixin-based inheritance. In *Proc. OOPSLA '90*, pages 303–311, 1990.
12. W. R. Cook. *A Denotational Semantics of Inheritance*. PhD thesis, Brown University, 1989.
13. E. Crank and M. Felleisen. Parameter-passing and the lambda calculus. In *Proc. POPL '91*, pages 233–244, 1991.
14. M. Felleisen and R. Hieb. The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103(2):235–271, 1992.
15. M. Flatt, S. Krishnamurthi, and M. Felleisen. Classes and mixins. In *Proc. POPL '98*, pages 171–183, 1998.
16. M. Flatt, S. Krishnamurthi, and M. Felleisen. A Programmer's Reduction Semantics for Classes and Mixins. In *Formal Syntax and Semantics of Java*, volume 1523 of *LNCS*, pages 241–269, 1999.
17. M. V. Limberghen and T. Mens. Encapsulation and composition as orthogonal operators on mixins: a solution to multiple inheritance problems. *Object Oriented Systems*, 3(1):1–30, 1996.
18. I. Mason and C. Talcott. Programming, transforming, and proving with function abstractions and memories. In *Proc. ICALP '89*, pages 574–588. LNCS 372, Springer-Verlag, 1989.
19. A. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.

A Type Rules

The type rules for class-related forms were presented in Section 4. The remaining type rules are presented here.

A.1 Subtyping Rules

The subtyping rules are standard. Objects support both depth and width subtyping.

$$\begin{array}{c}
\frac{}{\Gamma, \iota_1 <: \iota_2 \vdash \iota_1 <: \iota_2} \quad (<: \text{proj}) \qquad \frac{}{\Gamma \vdash \tau <: \tau} \quad (\text{refl}) \\
\\
\frac{\Gamma \vdash \tau_1 <: \tau_2 \quad \Gamma \vdash \tau_2 <: \tau_3}{\Gamma \vdash \tau_1 <: \tau_3} \quad (\text{trans}) \qquad \frac{\Gamma \vdash \tau' <: \tau \quad \Gamma \vdash \sigma <: \sigma'}{\Gamma \vdash \tau \rightarrow \sigma <: \tau' \rightarrow \sigma'} \quad (\text{arrow}) \\
\\
\frac{\Gamma \vdash \tau_i <: \sigma_i \quad i \in I \quad I \subseteq J}{\Gamma \vdash \{m_i : \tau_i\}^{i \in I} <: \{m_j : \sigma_j\}^{j \in J}} \quad (<: \text{record})
\end{array}$$

A.2 Type Rules for Expressions

The type rules for expressions other than class-related forms are simple, except for heaps, which have to be typed globally.

$$\frac{\text{typeof}(\text{const}) = \tau}{\Gamma \vdash \text{const} : \tau} \quad (\text{const}) \qquad \frac{}{\Gamma, x : \tau \vdash x : \tau} \quad (\text{proj}) \qquad \frac{\Gamma, x : \tau \vdash e : \sigma}{\Gamma \vdash \lambda x. e : \tau \rightarrow \sigma} \quad (\lambda)$$

$$\begin{array}{c}
\frac{\Gamma \vdash e_1 : \tau \rightarrow \sigma \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash e_1 e_2 : \sigma} \quad (app) \qquad \frac{}{\Gamma \vdash fix : (\sigma \rightarrow \sigma) \rightarrow \sigma} \quad (fix) \\
\\
\frac{\Gamma \vdash e : \tau \quad \Gamma \vdash \tau <: \sigma}{\Gamma \vdash e : \sigma} \quad (sub) \qquad \frac{\Gamma \vdash e_i : \tau_i}{\Gamma \vdash \{x_i = e_i\}^{i \in I} : \{x_i : \tau_i\}} \quad (record) \\
\\
\frac{\Gamma \vdash e : \{x : \sigma\}}{\Gamma \vdash e.x : \sigma} \quad (lookup) \qquad \frac{}{\Gamma \vdash ref : \tau \rightarrow \tau \text{ ref}} \quad (ref) \qquad \frac{}{\Gamma \vdash ! : \tau \text{ ref} \rightarrow \tau} \quad (!) \\
\\
\frac{}{\Gamma \vdash := : \tau \text{ ref} \rightarrow \tau \rightarrow \tau} \quad (:=) \\
\\
\frac{\Gamma' = \Gamma, x_1 : \tau_1 \text{ ref}, \dots, x_n : \tau_n \text{ ref} \quad \Gamma' \vdash v_i : \tau_i \quad \Gamma' \vdash e : \tau}{\Gamma \vdash H\langle x_1, v_1 \rangle \dots \langle x_n, v_n \rangle . e : \tau} \quad (heap)
\end{array}$$