

First-Class Object Sets

Erik Ernst

University of Aarhus, Denmark
eernst@cs.au.dk

Abstract. Typically, objects are monolithic entities with a fixed interface. To increase the flexibility in this area, this paper presents first-class object sets as a language construct. An object set offers an interface which is a disjoint union of the interfaces of its member objects. It may also be used for a special kind of method invocation involving multiple objects in a dynamic lookup process. With support for feature access and late-bound method calls object sets are similar to ordinary objects, only more flexible. The approach is made precise by means of a small calculus, and the soundness of its type system is shown by a mechanically checked proof in Coq.

Key words: Object sets, composition, multi-object method calls, types.

1 Introduction

In an object-oriented setting the main concept is the object. It is typically a monolithic entity with a fixed interface, such as a fixed set of messages that the object accepts. This paper presents a language design where sets of objects are first-class entities, equipped with operations that enable such object sets to work in a similar way as monolithic objects. An object set offers an interface to the environment which is a disjoint union of the interfaces of its members, and it supports cross-member operations, known as *object set method calls*, which are similar in nature to late-bound method calls on monolithic objects. The object set thus behaves in a way which resembles the behavior of a monolithic instance of a ‘large’ class that combines all the classes of the members of the object set, e.g., by mixin composition or multiple inheritance.

Object sets are useful because they are more flexible than such a monolithic instance of a large class: There are no restrictions on which classes may be put together in the creation of an object set, and there is no need to declare a large, composite class and refer to that class by name everywhere. In fact, object set types are structural with member class granularity—the type of an object set is a set of classes, and every subset is a supertype. Moreover, object sets could be modified during their life time, which would correspond to a change of class in the monolithic case.

On the other hand, access to a feature of an object set requires explicit selection of the member class which provides this feature, and the object set method call mechanism is quite simple rather than convenient. This is because

the emphasis in this language design has been put on expressing the required *primitives*, rather than giving the design of a fragment of a convenient and pragmatic programming language.

In fact, work on the implementation of the language gbeta [1–3] served as our starting point for the design of object sets. Objects in gbeta have a semantics which may be represented by collections of instances of mixins, i.e., as multi-entity phenomena rather than monolithic entities. Like object sets, they provide an interface which is a disjoint union of the interfaces of the included mixins, but unlike object sets there is no need to specify explicitly which mixin to use when accessing a feature. Like object sets, gbeta objects can have cross-entity features (such as methods or inner classes, which are then known as virtual), but unlike object sets these features are accessed in exactly the same way as single-mixin features. Similarly, since all gbeta objects are conceptually object sets there is no distinction (syntactically or otherwise) between the usage of object sets and “ordinary objects”.

The language gbeta also supports dynamic change of class for existing objects, and this corresponds to the replacement of the contents of the object by a larger object set. In context of the features included in this paper this operation is simple and safe, though of course it would require addition of mutable references to object sets to make it work as a dynamic change of class. In context of gbeta it is considerably more complex because dynamic specialization of an object may have effects that correspond to a dynamic replacement of actual type arguments of the class of the (multi-entity) object by some subtypes, which may cause a run-time error, e.g., because the value of an instance variable may thus become type incorrect. Because of this, dynamic object specialization in gbeta has been extended with restricted versions that are safe, but it is beyond the scope of this paper to model these refinements of the concept. Nevertheless, it is worth noting that it is possible to embody the primitives presented in this paper into a full-fledged programming language in such a way that they are convenient to use.

The contributions of this paper is the notion of object sets, the precise definition of their semantics and typing in a formal calculus, FJ_{set} , and the mechanically checked proof of soundness [4] for this calculus, using the Coq [5] proof assistant.

The rest of the paper is organized as follows: Section 2 presents the calculus informally and discusses the design. Next, Sect. 3 gives the formal definitions, and Sect. 4 describes the soundness result. Finally, Sect. 5 describes related work, and Sect. 6 concludes.

2 An Informal Look at the FJ_{set} Calculus

The FJ_{set} calculus is derived from the Featherweight Java calculus [6] by adding the object set related operations, allowing covariant method return types, and removing casts. The covariant method return types are included because they

are useful and standard today, and the casts are left out because they do not provide extra benefits in this context.

The crux of this calculus is of course the ability to express and use object sets. An object set is a set of objects collected into a single, typed entity. An object set may be decomposed in order to use individual members of the set, and used as a whole in a special kind of method call, the *object set method call*. Each object set is associated with a set of classes and each object in the set is uniquely associated with one particular class in the set of classes. Another way to describe this would be to say that each object in the set is *labeled* by a class. The object is an instance of that class or a subclass thereof. This makes it possible to access the object set members and to use each one of them according to an interface that it is known to support.

The correspondence between objects and classes in an object set is maintained by considering the set of objects and the set of classes as lists, and pairing up the lists element by element. This is possible for an object set creation expression (a variant of the well-known `new` expression for monolithic objects) because such an expression contains the two lists syntactically, and this ensures that every object set from its creation has a *built-in definition of the mapping* from classes to member objects. It also equips the members of the object set with an ordering. This ordering is insignificant with respect to typing, but it is significant with respect to the dynamic semantics, because it determines which method implementation is most specific during an object set method call. In other words, the ordering of the members of an object set is an implementation detail—crucial in the definition of its internal structure and behavior, but encapsulated and invisible at the level of types.

However, an expression denoting an object set may by subsumption be typed with an arbitrary subset of the associated classes, and they may be listed in an arbitrary order in the type. In other words, it is possible to forget some of the objects and also to ignore the ordering of the objects. However, the dynamic semantics only operates on an object set when it has been evaluated to such an extent that it is an object set creation at top level. This ensures that the object set operations are consistent because they are based on the built-in mapping.

Two operations are provided to decompose an object set. They both rely on addressing a specific member of the set via its associated class. One operation provides access to the object associated with the given class, and the other operation deletes that object and class from the object set, thus producing a smaller object set.

The object set method call operation is provided in order to gather contributions from all suitable objects in an object set, in a process that resembles a fold operation on a list. The call is based on an ordinary method whose signature must follow a particular pattern, namely that it takes a positive number of arguments and the type of the first argument is identical to the return type of the method. This makes it possible for the method to accept an arbitrary list of “ordinary” arguments—the arguments number two and up—and also to accept and return a value which plays the role as an accumulator of the final result. In this

```

1  class Printable extends Object {
2      String print(String s) { return "Plain Printable"; }
3      String separator() { return ", "; }
4  }
5  class Agent extends Printable {
6      String id;
7      String print(String s) { return s+" "+id; }
8  }
9  class Person extends Printable {
10     String name;
11     String print(String s) { return name+" "+s; }
12     String separator() { return " -- "; }
13 }
14 class Main extends {
15     { Printable } p;
16     String doPrint() {
17         return p.print@Agent("The name is") +
18                p@Printable.separator() +
19                p.print@Printable("");
20     }
21 }
22
23 // the following yields "The name is Bond -- James Bond"
24 new Main(new {Agent, Printable} (new Agent("Bond"),
25                                 new Person("James")))
26 .doPrint()

```

Fig. 1. An small example program in FJ_{set} .

sense the object set method call supports iteration over the selected members of the object set and collection of contributions to the final result, not unlike a folding operator applied to a list. Note that an object set method call does not require static knowledge about the type of any of the objects in the set.

The design of the object set method call mechanism was chosen to enable iteration over a subset of the members of the object set supporting a specific interface, without adding extra language mechanisms. Pragmatically, it might be more natural to use actual iteration in an imperative setting, or to return a data structure like an array containing the eligible object set members. But in this context we prefer a minimal design, and hence we ended up choosing the programmer convention driven approach based on ordinary nested method calls.

Figure 1 shows a small example program in FJ_{set} . This program shows how to create objects and object sets, how to perform an object set method call, how to decompose an object set in order to use a feature of one of its members, and it indicates the result of the computation. In order to make the example compact and readable it uses an extension of the calculus that includes a `String` type, string literals, and concatenation of strings with the `+` operator.

Lines 1–13 define three classes to support modeling a human being from two different points of view in an object set; the only difference from standard Java code is that there are no constructors, but the constructors in FJ style calculi are trivial and somewhat of an anomaly so we left them out. Note that the signature of the `print` method is such that it can be used for object set method calls: Its return type is also the type of the first (and only) argument.

The class `Main` has an instance variable (line 15) whose type is an object set, `{ Printable }`, which means that it is guaranteed that there is an object labeled `Printable` in this object set, but there may be other objects as well. The `doPrint` method (line 16–20) makes two object set method calls (line 17 and 19) and one ordinary method call (line 18), and returns the concatenation of the results. The object set method call on line 17 involves only one member of `p`, because only the first one is labeled by `Agent` or a subclass thereof. The call on line 19 involves both objects in `p`. The expression `p@Printable` on line 18 extracts the object labeled as `Printable` in `p`, which is the `Person` object, and calls its `separator` method, which by ordinary late binding returns " -- ".

Finally, note that subsumption makes it possible for the instance variable `p` to refer to an object set of type `{ Agent, Printable }`, and also that the usage of different classes in the object set method call can be used to filter the contributors to such a call in various ways.

3 The FJ_{set} Calculus

We now proceed to present the syntax, the dynamic semantics, and the type system of the FJ_{set} calculus, interspersed with short discussions about why the calculus is designed the way it is. We also give some remarks on how the presentation in this paper and the accompanying Coq proof fit together, based on the assumption that this kind of knowledge is useful for the development of a strong culture of using proof assistant software.

3.1 Syntax and Notation

A program is a class table and a main expression, and the semantics of a program is to execute the main expression in context of the given classes. As is common, we assume the existence of a fixed, globally accessible class table, `CT`, which lists all the class definitions in the program.

The syntax of the calculus from the level of classes and down is shown in Fig. 2. Notationally, we use overbars to denote lists of terms, so \bar{C} stands for the list $C_1 C_2 \dots C_n$ for some natural number n ; $n=0$ is allowed and yields the empty list, `'•'`. There may be separators such as commas or semicolons between the elements of such a list, but they are implicit and implied by the context.

Several constructs in the syntax are identical to the ones known from Featherweight Java. Class and method definitions are standard, using the variant of Featherweight Java that omits explicit constructors. The standard expressions are variables, field lookups, method calls, and `new` expressions.

Q	$::=$	<code>class C extends D { \overline{T} \overline{f}; \overline{M} }</code>	<i>class declarations</i>
M	$::=$	<code>T $m(\overline{T} \overline{x})$ { return e; }</code>	<i>method declarations</i>
e	$::=$	<code>x e.f e.m(\overline{e}) e@C e\C new C(\overline{e}) new {\overline{C}}(\overline{e}) e.m@C(\overline{e})</code>	<i>expressions</i>
v	$::=$	<code>new C(\overline{v}) new {\overline{C}}(\overline{v})</code>	<i>values</i>
T, U	$::=$	<code>C {\overline{C}}</code>	<i>types</i>
<code>Object, this</code>			<i>predefined names</i>
<code>C, D</code>			<i>class names</i>
<code>f, g</code>			<i>field names</i>
<code>x</code>			<i>variable names</i>
<code>m</code>			<i>method names</i>
<code>N</code>			<i>any kind of name</i>

Fig. 2. Syntax of FJ_{set} .

The remaining expressions are concerned with object sets. A class selection expression, $e@C$, provides the object labeled with the class C from the object set e . A class exclusion expression, $e\C$, provides an object set from which the object labeled with C as well as C itself has been deleted. The expression `new { \overline{C} }(\overline{e})` denotes creation of an object set which contains each of the objects denoted by the expression list \overline{e} , labeled by the list of classes \overline{C} .

Finally, the expression `e.m@C(\overline{e})` denotes an object set method call, which selects all objects from the object set e which are labeled with the class C or a subclass thereof, and calls a method m on each of them in the order they appear in the class list of the built-in mapping of the object set e . The method m must be defined in or inherited by the class C , and it must take a non-zero number of arguments where the first argument has the same type as the method return type, in order to enable the nested method call process mentioned in Sect. 1.

3.2 Auxiliary methods, Subtyping, and Wellformedness

Figure 3 defines the auxiliary functions used for field lookup and similar tasks. They are standard except for the function *distinct* which simply expresses that a given list of names (of any kind such as class names, method names, etc.) are distinct. As is common, quoting a class definition as a premise of a rule indicates the requirement that CT must contain that class definition.

The rules in Fig. 4 show subclassing ($\vdash C \sqsubset D$), which is standard, subtyping for object sets ($\vdash T \subset U$), which corresponds to the superset relation among the sets, and subtyping, which combines the two. Furthermore the judgement $C \in: \{\overline{C}\}$ holds whenever C is a superclass of one of the classes \overline{C} ; this is used in the dynamic semantics of object set method calls.

$$\begin{array}{c}
fields(\text{Object}) = \bullet \\
\\
\frac{\text{class } C \text{ extends } D \{ \overline{Uf}; \overline{M} \}}{\text{fields}(C) = \overline{Ug}, \overline{Tf}} \\
\\
\frac{\text{class } C \text{ extends } D \{ \overline{Uf}; \overline{M} \} \quad T m(\overline{T x}) \{ \text{return } e; \} \in \overline{M}}{mBody(m, C) = \overline{x.e}} \qquad \frac{\text{class } C \text{ extends } D \{ \overline{Uf}; \overline{M} \} \quad m \notin \overline{M} \quad mBody(m, D) = \overline{x.e}}{mBody(m, C) = \overline{x.e}} \\
\\
\frac{\text{class } C \text{ extends } D \{ \overline{Uf}; \overline{M} \} \quad m \notin \overline{M} \quad mType(m, D) = (\overline{T} \rightarrow T)}{mType(m, C) = (\overline{T} \rightarrow T)} \qquad \frac{\text{class } C \text{ extends } D \{ \overline{Uf}; \overline{M} \} \quad T m(\overline{T x}) \{ \text{return } e; \} \in \overline{M}}{mType(m, C) = (\overline{T} \rightarrow T)} \\
\\
distinct(\bullet) \qquad \frac{N \notin \overline{N} \quad distinct(\overline{N})}{distinct(N \overline{N})}
\end{array}$$

Fig. 3. Auxiliary functions for FJ_{set} .

In the Coq formalization of the calculus, transitivity for subclassing includes the requirement that the two pairs of classes are distinct, i.e., that $C \neq C''$ and $C' \neq C'$. An easy induction shows that each of the two definitions of subclassing is able to derive all the subclass judgements of the other. However, in order to show in Coq that subclassing is decidable, the addition of these requirements solves a problem because it is hard to specify in Coq that the derivation tree for a subtyping judgement is finite. As we shall see later on, the required finiteness is a consequence of the rule for class typing.

The requirement in the rule for object set subtyping that the classes \overline{D} are distinct is necessary in order to prevent occurrences of class lists with duplicate elements. It is only required for the supertype, $\{\overline{D}\}$, because distinctness for the subtype is ensured by other rules, in particular in the typing of object set creation expressions shown below in Fig. 7.

The type wellformedness requirements are shown in Fig. 5. They state that a class name is well-formed if there is a class of that name in the class table, and that an object set type, $\{\overline{C}\}$, is well-formed if it consists of distinct class names. Finally CT must satisfy $\text{Object} \notin CT$. Note that the class names in an object set type are not explicitly required to be defined in CT , because this requirement is a consequence of other rules. In general, the wellformedness requirements in this calculus are sufficient to enable the proof of soundness, but they are also minimal in the sense that removing any of them invalidates the proof. We believe that the use of proof assistant software may tighten the specification of well-formedness requirements in calculi, which is an area that otherwise easily gets a slightly imprecise treatment.

$$\begin{array}{c}
\frac{\text{class } C \text{ extends } D \{ \overline{Tf}; \overline{M} \}}{\vdash C \sqsubseteq D} \qquad \frac{\vdash C \sqsubseteq C'' \quad \vdash C'' \sqsubseteq C'}{\vdash C \sqsubseteq C'} \\
\\
\vdash C \sqsubseteq C \qquad \frac{\overline{D} \subseteq \overline{C} \quad \text{distinct}(\overline{D})}{\vdash \{\overline{C}\} \sqsubseteq \{\overline{D}\}} \\
\\
\frac{\vdash T \sqsubseteq U}{\vdash T < U} \qquad \frac{\vdash T \sqsubseteq U}{\vdash T < U} \qquad \frac{\vdash C_i \sqsubseteq C}{C \in \{\overline{C}\}}
\end{array}$$

Fig. 4. Subclassing and subtyping for FJ_{set} .

$$\begin{array}{c}
\vdash \text{Object OK} \qquad \frac{C \in \mathcal{CT}}{\vdash C \text{ OK}} \qquad \frac{\text{distinct}(\overline{C})}{\vdash \{\overline{C}\} \text{ OK}}
\end{array}$$

Fig. 5. Wellformedness rules for FJ_{set} .

3.3 Expression Evaluation

The dynamic semantics of FJ_{set} is presented in Fig. 6. Selection of a redex in a larger expression is defined in terms of evaluation contexts rather than congruence rules; they are listed at the bottom of the figure, where E denotes an expression with exactly one hole and E^* denotes a (non-empty) list of expressions with exactly one hole. With respect to the evaluation order, this calculus follows the tradition from FJ whereby the evaluation order is restricted as little as possible, and particular strategies like call-by-value are available as one of the possible choices.

The rules for field lookup and method invocation are standard. The rule for class selection, (R-SELECT), selects the member of the given object set labeled with the specified class. This rule serves as an example of the evaluation order issue mentioned above: evaluation has to proceed until the top level expression is an object set creation expression ($\text{new } \{\overline{C}\} (\dots)$) in order to reveal $\{\overline{C}\}$ and thus the built-in mapping of the object set, but the arguments need not be fully evaluated.

We stated earlier that an object set offers an interface which is a disjoint union of the interfaces of its members. The class selection operation fulfills this promise as follows: For a given object set, the classes used to label some members of the object set are made explicit in its type (others may have been lost by subsumption). The interface of the object set is the union of the interfaces of these classes, and thus the object set supports access to all these features of

$$\begin{array}{c}
\frac{\text{fields}(\mathbb{C}) = \overline{\mathbf{Tf}}}{\text{new } \mathbb{C}(\overline{\mathbf{e}}) . \mathbf{f}_i \rightsquigarrow \mathbf{e}_i} \\
\text{(R-FIELD)}
\end{array}
\qquad
\frac{}{\text{(new } \{\overline{\mathbb{C}}\} (\overline{\mathbf{e}})) \setminus \mathbb{C}_i \rightsquigarrow \text{new } \{\overline{\mathbb{C}} \setminus \#i\} (\overline{\mathbf{e}} \setminus \#i)} \\
\text{(R-DROP)}$$

$$\frac{\frac{\text{mBody}(\mathbf{m}, \mathbb{C}) = \overline{\mathbf{x}} . \mathbf{e}_0}{(\text{new } \mathbb{C}(\overline{\mathbf{e}})) . \mathbf{m}(\overline{\mathbf{e}'}) \rightsquigarrow} \quad [\text{this/new } \mathbb{C}(\overline{\mathbf{e}}), \overline{\mathbf{x}}/\overline{\mathbf{e}'}] \mathbf{e}_0}{\text{(R-INVK)}}$$

$$\frac{\begin{array}{l} \vdash \mathbb{C}_i \sqsubset: \mathbb{C} \quad i = \min\{j \mid \vdash \mathbb{C}_j \sqsubset: \mathbb{C}\} \\ \mathbf{v}_i = \text{new } \mathbb{D}(\overline{\mathbf{v}'}) \quad \text{mBody}(\mathbf{m}, \mathbb{D}) = \overline{\mathbf{x}} . \mathbf{e}_0 \\ \mathbf{e}'_0 = [\text{this/new } \mathbb{D}(\overline{\mathbf{v}'}) , \overline{\mathbf{x}}/\overline{\mathbf{e}'}] \mathbf{e}_0 \end{array}}{\frac{}{(\text{new } \{\overline{\mathbb{C}}\} (\overline{\mathbf{v}})) . \mathbf{m}@\mathbb{C}(\overline{\mathbf{e}}) \rightsquigarrow} \quad (\text{new } \{\overline{\mathbb{C}} \setminus \#i\} (\overline{\mathbf{v}} \setminus \#i)) . \mathbf{m}@\mathbb{C}(\overline{\mathbf{e}'_0})} \\
\text{(R-SINVK)}$$

$$\frac{}{(\text{new } \{\overline{\mathbb{C}}\} (\overline{\mathbf{e}})) @ \mathbb{C}_i \rightsquigarrow \mathbf{e}_i} \\
\text{(R-SELECT)}$$

$$\frac{\mathbb{C} \notin: \{\overline{\mathbb{C}}\}}{(\text{new } \{\overline{\mathbb{C}}\} (\overline{\mathbf{v}})) . \mathbf{m}@\mathbb{C}(\mathbf{e}, \overline{\mathbf{e}}) \rightsquigarrow \mathbf{e}} \\
\text{(R-SINVK-DONE)}$$

$$\begin{array}{l}
\mathbf{E} ::= \quad [_] \mid \mathbf{E} . \mathbf{f} \mid \mathbf{E} . \mathbf{m}(\overline{\mathbf{e}}) \mid \mathbf{e} . \mathbf{m}(\mathbf{E}^*) \mid \text{new } \mathbb{C}(\mathbf{E}^*) \mid \text{new } \{\overline{\mathbb{C}}\} (\mathbf{E}^*) \mid \mathbf{E} @ \mathbb{C} \mid \\
\mathbf{E} \setminus \mathbb{C} \mid \mathbf{E} . \mathbf{m}@\mathbb{C}(\overline{\mathbf{e}}) \mid \mathbf{e} . \mathbf{m}@\mathbb{C}(\mathbf{E}^*) \\
\mathbf{E}^* ::= \quad \overline{\mathbf{e}} \mathbf{E} \overline{\mathbf{e}'}
\end{array}$$

Fig. 6. Evaluation rules and evaluation contexts for FJ_{set} .

all those members. There are no naming conflicts because the choice of class is made explicit, i.e., it is a disjoint union. In a full-fledged language it is much more convenient to avoid the explicit class selection, but this is trivial in the cases where there is no naming conflict, and it should be handled explicitly when a conflict exists; the language *gbeta* uses this approach.

The rule for class exclusion, (R-DROP), deletes the requested class and the corresponding member from the object set. This rule introduces notation for a simple function that deletes the i 'th element from a list, namely $\overline{\mathbf{t}} \setminus \#i$, where $\overline{\mathbf{t}}$ are terms of any kind, e.g., class names or expressions. Usage of this notation implies that the list is long enough to contain the position to delete.

In the Coq formalization of this calculus the (R-DROP) rule zips the list of classes and the list of expressions together to a list of pairs, then deletes the pair which contains the specified class, and then unzips the shortened list of pairs to get the resulting list of classes and list of expressions. This is a relatively typical situation where the convenient formalization in Coq does not correspond exactly to a well-known or convenient notation for presentation in a paper, but the slightly awkward notation $\overline{\mathbf{t}} \setminus \#i$ used to express deletion-by-position seems to be the most readable way to bridge the gap.

The object set method call semantics is specified by two rules, (R-SINVK) and (R-SINVK-DONE). As mentioned, an object set method call amounts to a composite operation which includes a method call on each of the members of the set labeled by a class supporting that method. The rule (R-SINVK) specifies what to do when the object set contains a member supporting the requested method m , and the rule (R-SINVK-DONE) specifies what to do in the end when all such objects have been processed.

Whether an object set member supports m is determined by requiring that the member is associated with a subclass of the class C specified in the object set call. This means that each selected object will be an instance of C or one of its subclasses, and the method m will be defined for that object, with a signature which is identical to the signature of m in C , except for possible covariance in the return type. Objects supporting unrelated methods with the same name m are ignored.

The rule (R-SINVK) specifies how to call one method m and provide the results produced by this method call to the next method call. It requires that the list of classes \bar{C} associated with the object set contains a subclass C_i of the class requested in the call, C , and selects the ‘first’ one (the one with the smallest index i). It then removes the selected object from the receiver object set and repeats the object set method call with the result of the invocation of m on the selected object as its first argument. Note that the minimality of i is not needed for soundness, it is needed in order to ensure that object set method calls have a predictable semantics: it should accumulate the results from its members according to their built-in ordering.

However, the first argument does not look like a method call, it is actually given as $[\text{this/new } D(\bar{v}'), \bar{x}/\bar{e}\bar{e}]e_0$, but inspection of the rule for method call, (R-INVK), reveals that this is the result of taking one evaluation step after the method invocation $\text{new } D(\bar{v}') . m(\bar{e}\bar{e})$. It is necessary to express the rule in this form in order to maintain the property that all rules are compositional.

A similar investigation shows that the receiver of the object set method call after the step in (R-SINVK) is the result of taking one step after excluding the selected class C_i from the receiver object set before the step. Compositionality again forces the rule to take that step rather than expressing the result in terms of an explicit class exclusion operation.

The semantics of an object set method call may thus seem to be expressible in terms of other operations, but this is not the case because there is no way to select the class C_i appropriately without this operation. A primitive could be provided in order to make such a selection, but we have not found any which enables the same functionality without requiring strictly more static knowledge about the contents of object sets.

Finally, the rule (R-SINVK-DONE) yields the first argument of the object set method call in the situation where no object in the object set can be selected.

$$\begin{array}{c}
\frac{\Gamma(\mathbf{x}) = \mathbf{T}}{\Gamma \vdash \mathbf{x} : \mathbf{T}} \\
\text{(T-VAR)}
\end{array}
\qquad
\frac{\Gamma \vdash \mathbf{e} : \{\overline{\mathbf{C}}\} \quad \vdash \mathbf{C}_i \text{ OK}}{\Gamma \vdash \mathbf{e}@\mathbf{C}_i : \mathbf{C}_i} \\
\text{(T-SELECT)}$$

$$\frac{\Gamma \vdash \mathbf{e} : \mathbf{C} \quad \text{fields}(\mathbf{C}) = \overline{\mathbf{T}} \overline{\mathbf{f}}}{\Gamma \vdash \mathbf{e}.\mathbf{f}_i : \mathbf{T}_i} \\
\text{(T-FIELD)}$$

$$\frac{\Gamma \vdash \mathbf{e} : \{\overline{\mathbf{C}}\}}{\Gamma \vdash \mathbf{e} \setminus \mathbf{C}_i : \{\overline{\mathbf{C}} \setminus \#i\}} \\
\text{(T-DROP)}$$

$$\frac{\Gamma \vdash \mathbf{e} : \mathbf{C} \quad \text{mType}(\mathbf{m}, \mathbf{C}) = (\overline{\mathbf{T}} \rightarrow \mathbf{T})}{\Gamma \vdash \overline{\mathbf{e}} : \overline{\mathbf{U}} \quad \vdash \overline{\mathbf{U}} <: \overline{\mathbf{T}}} \\
\Gamma \vdash \mathbf{e}.\mathbf{m}(\overline{\mathbf{e}}) : \mathbf{T} \\
\text{(T-INVK)}$$

$$\frac{\text{distinct}(\overline{\mathbf{C}})}{\Gamma \vdash \overline{\mathbf{e}} : \overline{\mathbf{U}} \quad \vdash \overline{\mathbf{U}} <: \overline{\mathbf{C}}} \\
\Gamma \vdash \mathbf{new} \{\overline{\mathbf{C}}\}(\overline{\mathbf{e}}) : \{\overline{\mathbf{C}}\} \\
\text{(T-SNEW)}$$

$$\frac{\text{fields}(\mathbf{C}) = \overline{\mathbf{T}} \overline{\mathbf{f}}}{\Gamma \vdash \overline{\mathbf{e}} : \overline{\mathbf{U}} \quad \vdash \overline{\mathbf{U}} <: \overline{\mathbf{T}}} \\
\Gamma \vdash \mathbf{new} \mathbf{C}(\overline{\mathbf{e}}) : \mathbf{C} \\
\text{(T-NEW)}$$

$$\frac{\Gamma \vdash \mathbf{e} : \{\overline{\mathbf{C}}\} \quad \Gamma \vdash \overline{\mathbf{e}} : \mathbf{T} \overline{\mathbf{T}}}{\text{mType}(\mathbf{m}, \mathbf{C}) = (\mathbf{T} \overline{\mathbf{T}} \rightarrow \mathbf{T})} \\
\Gamma \vdash \mathbf{e}.\mathbf{m}@\mathbf{C}(\overline{\mathbf{e}}) : \mathbf{T} \\
\text{(T-SINVK)}$$

Fig. 7. Type rules for FJ_{set} .

3.4 Typing

The type rules for FJ_{set} are shown in Fig. 7. The rules for the typing of variables, field lookups, ordinary method invocations, and ordinary object creation are standard.

The rule (T-SELECT) specifies that the target must be typable as an object set containing the requested class, and the resulting type is then that class. It would be easy to change this rule and (R-SELECT) to select a subclass, i.e., to allow for the selection of a class \mathbf{C} as long as $\mathbf{C} \in: \{\overline{\mathbf{C}}\}$, but this could prevent the selection of a class \mathbf{C}' from an object set that is also associated with some subclass \mathbf{C}'' of \mathbf{C}' or make the operation ambiguous, and since there is no depth subtyping for object set types it would not enhance the expressive power or the flexibility of the language.

The rule (T-DROP) specifies that the target must be typable as an object set that includes the class to exclude, which is then removed from the type of the object set to produce the result type. For the same reasons as above it would not be useful to allow the requested class to be a superclass of the excluded class. The rule (T-SNEW) specifies the typing of object set creations. It simply requires that the classes used as labels are distinct and that each member has a subtype of its associated class. The rule (T-SINVK) specifies how to type object set method calls. It requires that the receiver is typable as an object set, but

$$\begin{array}{c}
\frac{\text{override}(\mathbf{m}, \mathbf{D}, \mathbf{T}, \bar{\mathbf{T}}) \quad \Gamma; \text{this}:\mathbf{C} \vdash \mathbf{e}:\mathbf{U} \quad \vdash \mathbf{U} <: \mathbf{T} \quad \vdash \mathbf{T}, \bar{\mathbf{T}} \text{ OK} \quad \text{distinct}(\bar{\mathbf{x}})}{\vdash \mathbf{T} \quad \mathbf{m}(\bar{\mathbf{T}}\bar{\mathbf{x}})\{ \text{return } \mathbf{e}; \} \text{ OK in } \mathbf{C}, \mathbf{D}} \\
\text{(T-METHOD)} \\
\\
\frac{\vdash \mathbf{D} \text{ OK} \quad \vdash \bar{\mathbf{T}} \text{ OK} \quad \vdash \bar{\mathbf{M}} \text{ OK in } \mathbf{C}, \mathbf{D} \quad \vdash \mathbf{D} \not<: \mathbf{C} \quad \vdash \mathbf{D} <: \mathbf{Object} \quad \text{distinct}(\text{fields}(\mathbf{D}) \bar{\mathbf{f}}) \quad \text{distinct}(\text{names}(\bar{\mathbf{M}}))}{\vdash \text{class } \mathbf{C} \text{ extends } \mathbf{D} \{ \bar{\mathbf{T}}\bar{\mathbf{f}}; \bar{\mathbf{M}} \} \text{ OK}} \\
\text{(T-CLASS)} \\
\\
\frac{\text{mType}(\mathbf{m}, \mathbf{D}) \text{ is undefined}}{\text{override}(\mathbf{m}, \mathbf{D}, \mathbf{T}, \bar{\mathbf{T}})} \quad \frac{\text{mType}(\mathbf{m}, \mathbf{D}) = (\bar{\mathbf{T}} \rightarrow \mathbf{T}') \quad \vdash \mathbf{T} <: \mathbf{T}'}{\text{override}(\mathbf{m}, \mathbf{D}, \mathbf{T}, \bar{\mathbf{T}})}
\end{array}$$

Fig. 8. Class and method typing for FJ_{set} .

does not require anything about the set of classes associated with this object set. On the other hand, the method \mathbf{m} must be defined or inherited in the class \mathbf{C} , it must take at least one argument, and the type of the first argument must be identical to the return type, which is also the type of the entire object set method call.

It would be very easy to change the (T-SINVK) rule to require $\mathbf{C} \in: \{\bar{\mathbf{C}}\}$ and adapt the soundness proof accordingly, which would guarantee that the object set method call would include at least one actual method call, but this is not required for soundness. Similarly, it would be easy to relax the rule such that the return type only has to be a subtype of the type of the first argument rather than being identical to it, but it would be hard to exploit this information unless the rule were modified to enforce that there is at least one actual method call. Even then, the accumulation of contributions from several members of the object set would have to start “from scratch” at each member, because the type of the first argument is fixed. Hence, these variations do not seem to be worthwhile.

Finally, Fig. 8 shows the rules for class and method typing, i.e., rules that apply type checking to the entire program. As opposed to the traditional treatment, these rules include all the requirements needed for programs to be well-formed—for instance in order to avoid cyclic inheritance graphs.

The rule (T-METHOD) specifies that a method \mathbf{m} defined in a class \mathbf{C} with superclass \mathbf{D} must correctly override any definitions of \mathbf{m} available in the superclass, it must have a body whose type is a subtype of the declared return type, it must have distinct argument names, and the specified types must be well-formed. The only non-standard element here is the requirement that argument names must be distinct.

The rules for *override* are given at the bottom of the figure; they are used to specify when a definition of a method \mathbf{m} with argument types $\bar{\mathbf{T}}$ and return type

T is correct in relation to definitions available in the superclass D . It is standard except that it allows for covariance in the return type, just like the Java language of today.

The rule (T-CLASS) specifies the standard requirements that the superclass D , all field types, and all methods must be well-formed. Moreover, the superclass cannot be a subclass of C itself, which prevents cycles in the inheritance graph; and the superclass must be a subclass of `Object`, which ensures that all inheritance chains are finite. This finiteness ensures that subclassing is decidable, which is used in the progress proof. Finally, there are distinctness requirements for field and method names.

All in all, this is not much more involved or verbose than the usual class and method typing rules, but it is complete in the sense that there are no additional (informal and maybe even implicit) well-formedness rules about programs to worry about. We think that it would be useful to make program well-formedness fully explicit, as we have done it here; it is, of course, a consequence of using proof assistant software, because the proofs cannot be completed unless such things are made precise and included in the specification.

4 Soundness

The FJ_{set} calculus is sound, which is shown via the standard progress and preservation results:

Theorem 1 (Progress). *If e is an expression typeable by $\emptyset \vdash e : T$ then either e is a value or there exists an expression e' such that $e \rightsquigarrow e'$.*

Theorem 2 (Preservation). *If the expression e in the environment Γ is typeable by $\Gamma \vdash e : T$ and it can take the step $e \rightsquigarrow e'$, then $\Gamma \vdash e' : U$ for some type U such that $\vdash U <: T$.*

A complete proof of these properties which has been mechanically checked by the proof assistant Coq is available for download [4]. It consists of approximately 6500 lines of Gallina code, divided into approximately 3500 lines specifically on the calculus, and approximately 3000 lines of standard language metatheory facilities from the Coq tutorial given at POPL 2008 [7].

5 Related Work

Dynamic languages like Self [8, 9] support a very general and flexible style of composite objects by means of parent slots and genuine delegation. Object sets are less flexible, but in return they are statically typed.

Object sets are similar to extensible records in some ways. For instance, Gaster and Jones [10] define polymorphic, extensible records and unions based on row variables, i.e., mappings from labels to types. With object sets, the associated classes work as labels and types combined; this reduces the flexibility

because there cannot be two labels with the same type, but given that object sets are intended to model composite objects it would correspond to repeated inheritance to have more than one member associated with the same class, and this would preclude a surface level syntax where class selection is implicit due to the name clashes. Object sets as presented here do not support extension; this is because we consider a ‘`lacks C`’ construct which promises that there is no class `C` in this object set to be unmanageable in real-world software development. On the other hand the extensible records do not have a late-bound operation that corresponds to our object set method calls, it only uses statically known components.

A well-established approach to extensible records is the Haskell `HList` library [11], where Kiselyov, Lämmel and Schupke use type level natural numbers and a number of layers on top of that to support type safe heterogeneous lists. Such lists are actually nested tuples, and the approach relies heavily on being able to use large type expressions which are inferred and never show up in the source code. If explicit typing is considered a valuable source of documentation then object sets are more manageable because they abstract away from the ordering of elements, and they may provide access to an arbitrary set of members without depending on the internal structure, i.e., the order of known members and the presence of unknown members.

6 Conclusion

We have presented the concept of object sets as a first class language construct which is capable of emulating the main features of traditional, monolithic objects: access to the disjoint union of the features of all object set members in the type, and support for a kind of method calls whereby the choice of methods to call is made dynamically, corresponding to feature access and method calls for ordinary objects. However, object sets are more flexible than ordinary objects, because they combine the features of several classes (like mixins or multiple inheritance, but without the name clashes), and they provide the machinery needed in order to support dynamic metamorphosis of object sets. The mechanism is useful in its own right, but it is likely to benefit from a pragmatic layer on top of the operations shown in this paper, because this makes the syntax more compact and convenient. The mechanisms of this paper might then provide good service as primitives on main-stream platforms such as .Net or JVM, which would make these platforms better at handling flexible object models in a type safe manner.

Acknowledgments The design of the object set method call mechanism owes some very useful insights to Kim Birkelund. The Coq proof was developed from a starting point created by Bruno de Fraine which was a proof of soundness for Featherweight Java without casts, which again used a number of files from the POPL 2008 Coq tutorial. This was very valuable material in the process of getting up to speed in Coq.

References

1. Ernst, E.: gbeta – A Language with Virtual Attributes, Block Structure, and Propagating, Dynamic Inheritance. PhD thesis, DEVISE, Department of Computer Science, University of Aarhus, Aarhus, Denmark (June 1999)
2. Ernst, E.: Higher-order hierarchies. In Cardelli, L., ed.: Proceedings ECOOP'03. LNCS 2743, Heidelberg, Germany, Springer-Verlag (July 2003) 303–329
3. Ernst, E., Ostermann, K., Cook, W.R.: A virtual class calculus. In: Proceedings POPL'06, Charleston, SC, USA, ACM (2006) 270–282
4. Ernst, E.: Coq proof of soundness for the FJ_{set} calculus (October 2008) <http://www.daimi.au.dk/~eernst/Sw65ab/objsetproof.tgz>.
5. Bertot, Y., Castéran, P.: Interactive Theorem Proving and Program Development. Coq'Art: The Calculus of Inductive Constructions. Texts in Theoretical Computer Science. Springer Verlag (2004)
6. Igarashi, A., Pierce, B., Wadler, P.: Featherweight Java: A minimal core calculus for Java and GJ. TOPLAS **23**(3) (May 2001) 396–459
7. Aydemir, B.: Using proof assistants for programming language research (January 2008) <http://www.cis.upenn.edu/~plclub/pop108-tutorial/>.
8. Ungar, D., Smith, R.B.: Self: The power of simplicity. In: Proceedings OOPSLA'87, Orlando, FL (October 1987) 227–242
9. Agesen, O., Bak, L., Chambers, C., Chang, B.W., Hölzle, U., Maloney, J., Smith, R.B., Ungar, D., Wolczko, M.: The Self 4.0 Programmer's Reference Manual. Sun Microsystems, Inc., Mountain View, CA (1995)
10. Gaster, B.R., Jones, M.P.: A polymorphic type system for extensible records and variants. Technical Report NOTTCS-TR-96-3, Department of Computer Science, University of Nottingham (November 1996)
11. Kiselyov, O., Lammel, R., Schupke, K.: Strongly typed heterogeneous collections. In: Haskell Workshop. (2004) 96–107