Sharing Classes Between Families:
Technical report

Xin Qi   Andrew C. Myers

{qixin,andru}@cs.cornell.edu

Abstract

Class sharing is a new language mechanism for building extensible software systems. Recent work has separately explored two different kinds of extensibility: first, family inheritance, in which an entire family of related classes can be inherited, and second, adaptation, in which existing objects are extended in place with new behavior and state. Class sharing integrates these two kinds of extensibility mechanisms. With little programmer effort, objects of one family can be used as members of another, while preserving relationships among objects. Therefore, a family of classes can be adapted in place with new functionality spanning multiple classes. Object graphs can evolve from one family to another, adding or removing functionality even at run time.

Several new mechanisms support this flexibility while ensuring type safety. Class sharing has been implemented as an extension to Java, and its utility for evolving and extending software is demonstrated with realistic systems.

1   Introduction

It has long been observed that much of the difficulty of building reliable software systems arises from their size and complexity [5]. Language mechanisms for modular and object-oriented programming have helped make large systems easier to build and to reason about. However, it remains challenging to extend, evolve, and update software systems. Incremental updates to software systems may require coordinated changes to state and behavior that span multiple software components. We posit that new mechanisms are needed to enable software to be reused and extended in a modular, scalable, safe way.

Inheritance does provide a modular way to extend existing code with new behavior, which helps explain why large software systems are often built using object-oriented languages such as Java or C++. But ordinary inheritance has two serious limitations. First, interacting classes can only be extended individually; inheritance does not support changes that span multiple classes. Second, new functionality cannot be added to objects of an existing class without modifying the class definition—which would not be modular.

These limitations have been addressed separately by two lines of research. In the first, the ability to extend classes as a group is provided by *family inheritance* mechanisms, which provide inheritance at the granularity of a *family* of related classes rather than on individual classes. Family inheritance mechanisms include virtual classes [25, 18, 10], mixin layers [41], variant path types [23], and nested inheritance [29, 31]. The second, separate line of research has explored language support for *adaptation*. Adaptation mechanisms allow the modular addition of functionality to an existing class, without any change to the original class definition. Objects of the class are augmented with new operations or state. Mechanisms for adaptation include open classes [11], expanders [50], classboxes [3], and aspect binders [27]. Typically, adaptation operates on individual classes; it does not support coordinated changes while preserving relationships among augmented objects.

This paper integrates these two distinct approaches to extensibility for the first time. The result is a powerful new capability: interacting objects from a family of classes can be adapted with new functionality,

while preserving the identities of objects and the relationships among them. A single operation can cause an entire data structure to be augmented in place with new behavior.

The key new idea is *class sharing*: different families can have classes that are shared between the families. An object that is an instance of the shared class in one family is also an instance of the corresponding class in any family that shares the class. Objects can then be viewed from either family. Class sharing supports *bidirectional adaptation*: not only can objects of a base family be adapted into a derived family, but those of the derived family can be adapted to the base family.

Class sharing has been realized in a language J&$_s$ (pronounced "jet-ess"), an extension of the Java language that adds class sharing to the language J& [31]. In J&$_s$, each object reference provides a *view* that dynamically determines the family and thus the behavior of the object when used through that reference. The view of an object also transitively affects the behavior of other objects reached via that object reference, because it determines which family to interpret those objects with respect to. An explicit *view change* produces a reference with a new view, while preserving object identity.

Class sharing enables new ways to extend software systems:

- *Family adaptation.* New functionality spanning multiple classes can be added to an entire family in-place, adapting objects with new state and behavior.

- *Dynamic object evolution.*

  The family from which an object is viewed is determined dynamically by its reference, so collections of interacting objects can be updated with new behavior at run time, without changing the interface to existing clients. This allows running servers to be evolved with new state and behavior.

- *In-place translation.* In-place translation is the transfer of data structures from one family to another without creating new objects for shared classes and without side effects. We have used this approach to build compilers: an abstract syntax tree from one language can be translated to another language while updating only the nodes that require it.

Despite all this expressive power, J&$_s$ is type-safe—programs never create run-time errors—and type checking is modular. Because the behavior of an object depends on the family from which it is being viewed, designing a sound type system is challenging. To accomplish this, J&$_s$ uses *view-dependent types* to ensure that late-bound type names belong to consistent families, and uses *masked types* [36] to control access to fields that are not shared between families. The key features of J&$_s$ have been formalized in an object calculus, and the type system has been proved sound.

The rest of the paper proceeds as follows. Section 2 describes how class sharing works in the new language J&$_s$, and gives examples of its use. Section 3 shows how to safely support unshared state within shared classes. The formalization and proof of soundness is shown in Section 4 and Section 5. The implementation of J&$_s$ is described in Section 6. Section 7 shows that the performance of J&$_s$ is reasonable and describes its use to build realistic systems. Section 8 discusses related work, and Section 9 concludes.

2 Sharing classes

The new class sharing mechanism in J&$_s$ is a safe, modular mechanism that relaxes the disjointness of class families. A family of classes may not only inherit another family (and hence all its nested classes), but may also *share* some of the classes from the family it inherits from. This enables new kinds of extensibility.

```
class AST {                          class TreeDisplay {
  class Exp {...}                       class Node {
  class Value                             void display() {...}
     extends Exp {...}                  }
  class Binary                         class Composite extends Node {
     extends Exp                         Node getChild(int i) {...}
  { Exp l, r; }                        }
  ...                                  class Leaf extends Node {...}
}                                    }
```

Figure 1: An expression family and a GUI family

```
1  class ASTDisplay extends AST & TreeDisplay {
2     class Exp extends Node { void display() { ... } }
3     class Value extends Exp & Leaf { ... }
4     class Binary extends Exp & Composite
5        { void display() { ... l.display(); ... } }
6     void show(Exp e) { e.display(); }
7  }
```

Figure 2: Mixing display into expressions, with nested inheritance

2.1 Family inheritance

J&$_s$ builds on the family inheritance mechanisms introduced by Nystrom et al.: nested inheritance [29] and nested intersection [31].

Nested inheritance is inheritance at the granularity of a *namespace* (a package or a class), which defines a family in which related classes are grouped. Nested inheritance supports coordinated changes that span the entire family. When a namespace inherits from another (base) namespace, not only are fields and methods inherited, but also namespaces nested in the base namespace. In addition, the derived namespace can *override*, or *further bind* [26] inherited namespaces, changing nested class declarations. Nested inheritance does not provide adaptation, however. Nested classes in the inheriting namespace, even those not overridden (*implicit classes*), are different classes than those of the same name in the base namespace.

Nested intersection supports *composing* families with generalized *intersection types* [39, 12]. Given classes *S* and *T*, their intersection *S&T* inherits all members of *S* and *T*. When two namespaces are intersected, their common nested namespaces are themselves intersected, i.e., $(S\&T).C = (S.C)\&(T.C)$.

Figure 1 shows an example with two families of classes. Class AST contains a family of classes for representing expressions, and class TreeDisplay contains a family of classes for graphically visualizing trees. Figure 2 shows the skeleton of code that implements the functionality of displaying an expression as a tree—without changing the existing families. ASTDisplay inherits both AST and TreeDisplay, and therefore inherits all nested classes from both of them. The GUI classes are implicit in ASTDisplay, and expression classes are further bound to inherit GUI classes in addition to their original superclasses, and to override GUI methods with appropriate rendering code. As the show method in Figure 2 demonstrates, expression classes in the new family support the added display method.

Two mechanisms are essential for nested inheritance to work: *late binding of type names* and *exact types*. These mechanisms are also important for class sharing.

Late binding of type names. Late binding of type names ensures relationships between classes are preserved in the derived family.

When the name of a class is inherited into a new namespace, the name is interpreted in the context of the new namespace. For example, inside the family ASTDisplay, the type Exp refers to Exp in ASTDisplay.

3

Consider the field `l`, which is declared in `AST.Binary` with type `Exp`, and then inherited by the class `ASTDisplay.Binary`. When inherited, its type is the `Exp` of `ASTDisplay`, *not* the original type. This late binding makes the call `l.display()` on line 5 legal. Similarly, the superclass of `ASTDisplay.Binary` is `ASTDisplay.Exp`, not `AST.Exp`.

Two mechanisms make the late binding of type names type-safe: *dependent classes* and *prefix types*. The dependent class *p*.`class` represents the run-time class of the object referred to by the *final access path p*. A final access path is either a final local variable, including `this` and final formal parameters, or a field access *p*.`f`, where *p* is a final access path and `f` is a final field. In general, the class represented by *p*.`class` is statically unknown, but fixed. A prefix type *P*[*T*] represents the enclosing namespace of the class or interface *T* that is a subtype of the namespace *P*, i.e., the family at the level of *P* that contains *T* [29]. In Figure 2, if one writes `AST[this.class]`, it refers to either `AST` or `ASTDisplay`, depending on the run-time class of the value stored in `this`.

Type names that are not fully qualified are sugar for members of prefix types that depend on the current class `this.class`, and in an inheriting family, they will be reinterpreted. In Figure 1, the type `Exp` inside class `AST` is sugar for `AST[this.class].Exp`.

Exact types.   Both dependent classes and prefix types of dependent classes are *exact types* [6]: all instances of these types must have the same run-time class.

Simple types may be exact too. If *A* is a class, the exact type *A*! represents values of the run-time class *A*. Even if class *B* inherits from *A*, neither *B* nor *B*! is a subtype of *A*!. Exactness applies to the entire type preceding "!", so supertypes of a simple exact type can be obtained by shifting the exactness outward. For example, `ASTDisplay.Exp!` is not a subtype of `AST.Exp!`, but it is a subtype of `ASTDisplay!.Exp`, which is a subtype of `ASTDisplay.Exp`.

Exact types also restrict the subtyping relationships of nested types.   For example, `ASTDisplay!.Binary` is not a subtype of `AST!.Binary`, even though `ASTDisplay.Binary` is a subtype of `AST.Binary`. Therefore exact types can mark the boundary of a family: classes nested in `ASTDisplay!` form one family and those nested in `AST!` form another. Non-dependent exact types provide a *locally closed world*: at compile time, one can enumerate *all* classes that are subtypes of *A*!.*C*, without inspecting subclasses of *A*. Exact types are important for the modularity of J&*s* (Section 2.5).

2.2   Sharing declarations

J&*s* introduces shared classes with *sharing declarations* in the derived family, such as the declaration "`shares A.C`" in this code:

```
class A { class C ... } // the base family
class B extends A {     // the derived family
  class C extends D shares A.C { ... }
}
```

This sharing declaration establishes a *sharing relationship* between classes `A.C` and `B.C`. Sharing declarations induce an equivalence relation on classes that is the reflexive, symmetric, and transitive closure of the declared sharing relationships. If two classes have a sharing relationship, they have the same set of object instances. However, subclasses of the two shared classes are not automatically shared, unless the subclasses also have appropriate sharing declarations. Therefore, the sharing relationship established in the above example can be represented as a relationship between two exact types, written `A.C!` ↔ `B.C!`.

J&*s* requires that only an overriding class in a derived family (e.g., `B.C`) may declare a sharing relationship with the overridden class in a base family (e.g., `A.C`). This restriction helps keep shared classes similar to each other.

```
1  class ASTDisplay extends AST & TreeDisplay {
2    class Exp extends ... shares AST.Exp { ... }
3    class Value extends ... shares AST.Value { ... }
4    class Binary extends ... shares AST.Binary
5      { void display() { ... l.display(); ... } }
6    void show(AST!.Exp e) sharing AST!.Exp = Exp {
7      Exp temp = (view Exp)e;
8      temp.display();
9    }
10 }
```

Figure 3: Using class sharing to adapt `Exp` to `TreeDisplay`

Sharing vs. subtyping. The J&$_s$ language keeps subtyping largely separate from sharing. Adding a sharing relationship does not change the subtyping relation. In the above example, `B.C` is a subtype of `A.C`, and `B!.C` is not a subtype of `A!.C` due to the exactness, whether there is a sharing declaration or not. The sharing relationship does not make `A.C` a subtype of `B.C`, nor does it create any subtyping relationship between `A!.C` and `B!.C`.

Since sharing is not subtyping, it is in general not allowed to directly treat an object of a shared class as an instance of the other class in the sharing relationship; an explicit *view change* may be required (see Section 2.3 for more details).

Family adaptation. Sharing solves the problem of *object adaptation* [50], in which the goal is to augment existing objects with new behavior or state. Adaptation is different from inheritance, where only objects of new classes have the new behaviors. J&$_s$ is the first language to fully integrate family inheritance with in-place adaptation that preserves object identity. Moreover, because sharing is an equivalence relation on classes, J&$_s$ supports bidirectional, transitive adaptation.

Adaptation can improve the example code from Figures 1 and 2. Although `ASTDisplay` in Figure 2 provides expression classes extended with the ability to display themselves, this new functionality is not available for instances of the original classes in `AST`. This is unfortunate, because instances of the original classes might be created by existing library code or deserialized from a file. We can avoid this limitation by using class sharing as shown in Figure 3. Here, `shares` clauses cause the two families `ASTDisplay` and `AST` to share all the expression classes. Instances of `AST` expression classes are also instances of corresponding `ASTDisplay` expression classes.

Because of sharing, expression objects from the `AST` family can possess GUI display operations, even if the implementer of `AST` was not aware of `TreeDisplay` or `ASTDisplay`; client code—for example, a visualization toolkit, which expects `TreeDisplay` objects—would obtain the ability to handle existing `AST` objects. Thus every expression class in the `ASTDisplay` family becomes an *adapter* [19] for the corresponding class in `AST`; the `ASTDisplay` family provides *family adaptation* for `AST`.

Given a tree of expression objects from `AST`, family adaptation ensures that the whole tree is safely adapted to `ASTDisplay`. The adaptation preserves the original tree structure, and the relationships between objects in the tree. This contrasts with prior adaptation mechanisms that work on individual objects, which do not guarantee safety or the preservation of object relationships. With prior mechanisms such as the adapter design pattern [19], one might forget to adapt the left child of a `Binary` object, and the call `l.display()` on line 5 would fail.

In Figure 3, every expression class has a sharing declaration, which could be tedious to write. J&$_s$ provides the `adapts` clause as a shorthand for adding sharing declarations to all inherited member classes. For example, `ASTDisplay` may be declared with the following class header, without any individual sharing

declarations:

```
class ASTDisplay extends ... adapts AST { ... }
```

2.3 Views and view changes

Views. If two classes are shared, a single object might be treated as an instance of either one. Each class is a distinct *view* of that object. A J&$_s$ object can have any number of views, all equally valid. This contrasts with an ordinary object-oriented language like Java (or even J&), where an object has exactly one view: its run-time class.

In J&$_s$, an object reference is not just a heap location; it is essentially a pair $\langle \ell, T \rangle$ of a heap location ℓ and a type T, where T is the view, represented as a non-dependent exact type. The view T determines the behavior of the object when accessed through that reference.

For example, the method `display` in Figure 3 cannot be directly called on an object created as an instance of the class `AST.Binary`, because the reference has the view `AST.Binary!`. However, when the object obtains a new reference—for example, by storing the object in a local variable—with the view `ASTDisplay.Binary!`, it also obtains a new behavior—the method `display` becomes available through the new reference. Moreover, methods that are available through the original reference might behave differently when they are called through the new reference. See Section 2.4 for an example.

View changes. J&$_s$ has a *view change* operation $(\texttt{view}\ T)e$, which generates a new reference with the same heap location but a different view. On line 7 in Figure 3, the method `show` contains a view change expression $(\texttt{view Exp})e$, the result of which is a reference that still points to the same object as `e` but with a new view that is a subtype of `Exp`. (Recall from Section 2.1 that within `ASTDisplay`, `Exp` is sugar for `ASTDisplay[this.class].Exp`).

View changes support late binding. Although the expression $(\texttt{view}\ T)e$ has a statically known type T, the actual run-time view of the result is a subtype of T that is shared with the run-time view of the value of e. For example, in line 7 of Figure 3, if `e` evaluates to a value with the view `AST.Value!`, then within the family `ASTDisplay`, the expression $(\texttt{view Exp})e$ produces a value with the view `ASTDisplay.Value!`, which is a subtype of `ASTDisplay!.Exp`, and shared with `AST.Value`.

A view change $(\texttt{view}\ T)e$ looks syntactically like a type cast $(T)e$, and it does have the same static target type T, but it is actually quite different. First, a type cast might fail at run time, but view changes that type-check always succeed. Second, no matter whether it is an upcast or a downcast, an ordinary cast only works if the target type is a supertype of the *run-time* class of the object. However, a view change has in general a target type that is neither a supertype nor a subtype of the current view of the object, but from another family. For example, if families `A` and `B` share the class `C`, together with all its subclasses within each family, the following code is legal:

```
A!.C a = new A.C();
B!.C b = (view B!.C)a;
```

The initial view of the created object is `A.C!`. The target type `B!.C` in the view change is neither a supertype nor a subtype of `A.C!`. Of course, if the type in a view change is indeed a supertype of the current view of the instance, the view change is a no-op.

Third, a type cast checks run-time typing information, but has no other effects at run time. By contrast, a view change can affect the behavior of the object.

View-dependent types. Nested inheritance uses the dependent class p.`class` to indicate the family that the object referenced by p belongs to. J&$_s$ generalizes p.`class` to be *view-dependent*, since an object may be a member of multiple families. For example, suppose `ASTDisplay` contained the following code:

```
AST!.Binary a = new AST.Binary();
Binary b = (view Binary)a;
```

At run time, `a.class` would denote `AST.Binary!`, and `b.class` would denote `ASTDisplay.Binary!`. This example shows that in J&$_s$, the dependent classes associated with different aliases (`a` and `b` in this example) are not necessarily equal; they are interpreted as the views associated with the respective references.

Prefix types of dependent classes are also view-dependent, ensuring that late-bound type names belong to consistent families. Consider the left child of class `Binary`, stored in the field `l`. The type of the field is `Exp`, sugar for `AST[this.class].Exp`, which depends on the view of the object that contains the field. Accessing the left child with `a.l` returns an object in the family `AST`, whereas accessing it with `b.l` returns the same object, but with a view in the family `ASTDisplay`. In either case, the left child object and the containing `Binary` object have views that are in the same family. This means that for a tree of objects in one family, a single explicit view change on the root object effectively moves the whole tree to another family, implicitly triggering view changes on child objects as they are accessed.

2.4 Dynamic object evolution via view change

View-based dynamic method dispatching. In J&$_s$, method calls are dispatched on the current *view* of the receiver object, rather than on the receiver itself as in Java. When two references to the same receiver object have different views, the same method call may invoke different code. Therefore, when a J&$_s$ class overrides a method inherited from its declared shared class, both versions of the method become available to the object, and the choice is made at run time, based on the view associated with the reference.

Method dispatching in J&$_s$ differs from that of nonvirtual methods in C++ [44], and from the statically scoped adaptation in *expanders* [50]. For nonvirtual methods in C++, the static type of the receiver acts as a static view that selects the method to call. By contrast, the view change operation in J&$_s$ affects method dispatching, but still allows late binding, since the type T in a view change $(\text{view } T)e$ is a supertype of the statically unknown run-time view.

Expanders dispatch methods within the same family dynamically, but unlike with class sharing, the choice between original code and expander code is static. Expander methods can be invoked only when expanders are in scope, and therefore the behavior of existing code written before the expanders cannot change. Since expansion is *not* inheritance, expander methods can only *overload* original methods rather than *overriding* them.

Dynamic object evolution. View-based method dispatching enables a new form of dynamic object evolution, in which existing objects are updated with different behavior without breaking running code. It is more powerful than object adaptation, which generally only adds new behavior to existing objects, whereas evolution can also make objects change behavior, even in a context that does not mention the updated classes. J&$_s$ supports evolution at the family level, evolving interacting objects consistently to the new family.

For example, Figure 4 shows a package `service` that implements several network services, and a dispatcher that calls different services based on the kind of the received packet. The server code has a static field storing the dispatcher, and an event loop:

```
static service.Dispatcher disp;
...
while (true) { ... disp.dispatch(p); ... }
```

Suppose the system implemented in the `service` family has started running, and then an updated package `logService` is developed that extends `service` with logging at various places (Figure 4 only shows the additional logging in the `dispatch` method). The goal is to update the system with logging ability, without having to stop it from running.

```
package service;                          package logService extends
class SomeService {                          service;
  void handle(Packet p) {...}            class SomeService shares
}                                            service.SomeService {...}
...                                       ...
class Dispatcher {                        class Logger {...}
  SomeService s;                          class Dispatcher shares
  void dispatch(Packet p) {                service.Dispatcher {
    switch (p.kind) {                        Logger logger;
      case 0: s.handle(p);                   void dispatch(Packet p) {
      ...                                       logger.log(...);
    }                                          ...
  }                                          }
}                                          }
```

Figure 4: Evolution of a network service package

To evolve the system from `service` to `logService`, the view of the dispatcher object stored in the static field `disp` needs to be changed. This could be done in initialization code in the extended package, as follows:

```
service!.Dispatcher d =
  (service!.Dispatcher)Server.disp; // cast
Server.disp = (view Dispatcher)d;   // view change
```

After this view change, the `dispatch` method overridden in the extended package will be called.

More importantly, although just a single explicit view change is applied to the dispatcher object, all the other objects transitively reachable from the dispatcher, such as through the field `s` of type `SomeService`, will also obtain new views when they are accessed, resulting in using the versions of their methods that have logging enabled. Thus, a single explicit view change causes a consistent evolution of many objects to the new family. This kind of evolution is simpler to implement and likely to be more efficient than going through all the objects and updating them individually.

2.5 Sharing constraints

Sharing relationships between types are not always preserved by derived families, either by design or as the result of changed class hierarchy in the derived family. For example, in a family inherited from `ASTDisplay`, one might choose not to share any class, or to share classes from `TreeDisplay`, and in either case, the new family no longer shares classes with the `AST` family.

Therefore, a view-change operation that works in the base family might not make sense in the derived family. J&$_s$ does not try to check all inherited method code for inapplicable view changes, but rather makes checking modular via *sharing constraints*. A J&$_s$ method can have sharing constraints of the form sharing $T_1 = T_2$, which means that any value of type T_1 can be viewed as of type T_2, and vice versa, or in other words, that the method body may assume the sharing relationship $T_1 \leftrightarrow T_2$. A view change can only appear in a method with an enabling sharing constraint. Outside the scope of sharing constraints, the type checker (and programmer) need not be concerned with sharing. Therefore, reasoning about class sharing is local.

For example, in Figure 3, the method `show` has a sharing constraint `AST!.Exp = Exp` (line 6), which allows the view change `(view Exp)e` to be applied to the variable `e` of static type `AST!.Exp`.

To know statically that the view change `(view Exp)e` will succeed, we must know that every subclass of `AST!.Exp` has a corresponding shared subclass under `ASTDisplay!.Exp`. J&$_s$ requires that some prefix of

```
class A1 {
  class B { }
  class C { D g; }
  class D { }
}
class A2 extends A1 {
  class B shares A1.B {
    T f;                      // a new field
  }
  class C shares A1.C\g { } // shared with a mask
  class E extends D { }    // a new subclass of D
}
```

Figure 5: Shared classes with unshared fields

each type in the constraint is exact, and either non-dependent or only dependent on the path `this`; thus, we can check all the subclasses in a locally closed world (Section 2.1), without a whole-program analysis. In this example, the exact prefixes in question are `AST!` and `ASTDisplay[this.class]`.

The type checker verifies that sharing constraints in the base family still hold in the derived family; base family methods whose sharing constraints do not hold must be overridden.

Although sharing constraints support modular type checking, they do introduce an annotation burden for the programmer. Our experience suggests that the annotation burden is manageable. While it appears possible to automatically infer sharing constraints, by inspecting the type of the source expression and the target type of every view change operation in the method body, we leave this to future work.

3 Protecting unshared state

3.1 Unshared fields

In J&$_s$, shared classes do not necessarily share all their fields. Unshared fields are important for greater extensibility, but they pose challenges for the safety of the language. J&$_s$ uses *masked types* [36] and duplicate fields to ensure safety in the presence of unshared fields.

Figure 5 illustrates the two kinds of unshared fields. First, new object fields may be introduced by shared classes in the derived family. Because these fields do not exist in the base family, they cannot be shared. In Figure 5, classes `A1.B` and `A2.B` are shared, but `A2.B` introduces a new field `f` that does not exist in `A1.B`. When a view change from the base family to the derived family—for example, from `A1!.B` to `A2!.B`—is applied, the new field (e.g., the field `f` in Figure 5) would be uninitialized, which may be unexpected to the code in the derived family.

J&$_s$ uses masked types to prevent the possibly uninitialized new field from being read. A masked type, written $T\backslash f$, is the type T without read access to the field `f`. We say that the field `f` is *masked* in $T\backslash f$. The mask on a field can be removed with an assignment to that field. Masked types introduce the subtyping relationship $T \leq T\backslash f$. For the example in Figure 5, a view change from `A1!.B` to `A2!.B` must have a mask on the target type:

```
  A1!.B b1 = new A1.B();
  A2!.B\f b2 = (view A2!.B\f)b1;
```

Therefore, after this view change, the field `f` of `b2` cannot be read before it is initialized.

The second kind of unshared field is those with unshared types. In Figure 5, the field `g` has type `D`, and it cannot be shared, because in the `A2` family, `g` might store an object of class `E`. When a view change is applied

9

```
package base;                          package pair extends base;
abstract class Exp                     abstract class Exp
{ ... }                                  { abstract base!.Exp translate
class Var extends Exp {                     (Translator v); }
  String x;                            class Pair extends Exp
  ...                                    { Exp fst, snd; ...}
}                                      class Translator {
class Abs extends Exp {                  base!.Abs reconstructAbs
  String x;                               (Abs old, String x,
  Exp e;                                   base!.Exp exp) { ... }
  ...                                    ...
}                                      }
```

Figure 6: Lambda calculus and pair compiler extension

from the derived family to the base family (for example, from `A2!.C` to `A1!.C`), an `A2.E` object stored in g would not have a view compatible with the base family.

Therefore, J&$_s$ requires that fields with unshared types are masked in the sharing declaration. A duplicate field with the same name for the shared class is also generated automatically in the derived family. For example, in Figure 5, it is as if the class `A2.C` has its own implicit declaration of field g:

```
class C shares A1.C\g {
  D g;
}
```

An instance of `A1.C` and `A2.C` contains two copies of the field g, each appearing as a "new" field to the other class. Which copy is accessed depends on the current view of the instance through which the field g is accessed. Field duplication prevents objects of an unshared class from being accidentally accessed in other families. Therefore, interacting objects always have consistent views.

When shared classes have duplicate fields, it is up to the programmer how to keep the copies in sync, or even whether to keep them in sync. The programmer may choose to construct a corresponding object in the target family, storing it in the duplicate field, as in the example in Section 3.2, or just to leave the field masked in the target family.

3.2 In-place translation with unshared classes

J&$_s$ supports in-place translation of data structures between families in which not all the classes are shared (translation is trivial if all classes are shared). As the data structure is translated, some objects in the structure may remain the same, with only a view change, and other objects, particularly those of unshared types, are explicitly translated. One use of this kind of translation is for compilers, where the data structure is an abstract syntax tree, and many parts of the AST do not need to change during a given compiler pass.

Figure 6 shows the skeleton of a simple compiler implemented with family inheritance but without class sharing, modeled on the Polyglot extensible compiler framework [30]. This example translates the λ-calculus extended with pairs, into the ordinary λ-calculus. The package `base` defines the target language through class declarations for AST nodes of the λ-calculus (e.g., `Abs` for λ abstractions); the package `pair` extends the source language with one additional AST node `Pair`. Classes inherited from `base` are further bound in `pair` with `translate` methods that recursively translate an AST from `pair` to `base`. The class `Translator` provides the methods AST classes use to (re)construct nodes in the target family using translated versions of their child nodes; the method `reconstructAbs` does this for λ abstractions. However, without sharing, the translation from `pair` to `base` has to recreate the whole AST, even for trivial cases like `Var`, because the two families `base` and `pair` are completely disjoint despite their structural similarity.

10

```
1   package pair extends base;
2   abstract class Exp shares base.Exp {
3     abstract base!.Exp translate(Translator v);
4   }
5   class Var extends Exp shares base.Var { ... }
6   class Abs extends Exp shares base.Abs\e {
7     base!.Exp translate(Translator v) {
8       base!.Exp exp = e.translate(v);
9       return v.reconstructAbs(this, x, exp);
10    }
11  }
12  class Pair extends Exp {
13    Exp fst, snd;
14    base!.Exp translate(Translator v) {
15      return new ...; //  (λx.λy.λf.f x y) ⟦fst⟧ ⟦snd⟧
16    }
17  }
18  class Translator {
19    base!.Abs reconstructAbs(Abs old, String x,
20                               base!.Exp exp)
21    sharing Abs\e = base!.Abs\e {
22      if (old.x == x && old.e == exp) {
23        base!.Abs\e temp = (view base!.Abs\e)old;
24        temp.e = exp;
25        return temp;
26      }
27      else return new base.Abs(x, exp);
28    }
29    ...
30  }
```

Figure 7: In-place translation of the pair language

By contrast, Figure 7 shows the J&$_s$ code that does in-place translation. The base family remains the same as in Figure 6, and is omitted. Classes Var and Abs in pair are declared to share corresponding classes in base (lines 5–6). Pair is not shared, because it does not exist in base.

Consider the two shared classes pair.Abs and base.Abs. The type of their field e, which is base[this.class].Exp, is interpreted as pair!.Exp and base!.Exp in the two families. The two interpreted field types are not shared, because a value of type pair!.Exp might have run-time class pair.Pair, which has no corresponding shared class in the base family. Therefore the two Abs classes each have their own copies of the field e, and the sharing declaration on line 6 has a mask on e.

Similarly, the sharing constraint on line 21 also has masks, and so does the view change operation on line 23, where the Abs instance is reused if all its subexpressions have been translated in place. The sharing constraint, together with the corresponding view change operation, has a mask on the field e, in case it points to an instance of an unshared subclass of Exp, such as Pair. The mask is removed on line 24, after which the type of the variable temp becomes base!.Abs.

As shown in the above example, J&$_s$ uses masked types to prevent objects of unshared types, such as Pair, from being leaked into an incompatible family (here, base). Masked types do introduce some annotation burden. However, class sharing is expected to be used in practice between families that are similar to each other, where extensibility and reuse are needed and make sense. In that case, there should

programs	Pr	$::= \langle \overline{L}, e \rangle$
class declarations	L	$::= \texttt{class } C \texttt{ extends } T_1 \texttt{ shares } T_2 \; \{\overline{L}\,\overline{F}\,\overline{M}\}$
field declarations	F	$::= [\texttt{final}]\; T\, f = e$
method declarations	M	$::= T\, m(\overline{T}\,\overline{x}) \texttt{ sharing } \overline{Q} \; \{e\}$
sharing constraints	Q	$::= T_1 \rightsquigarrow T_2$
pure types	PT	$::= \circ \mid PT.C \mid p.\texttt{class} \mid P[PT] \mid \&\overline{PT} \mid PT!$
types	T	$::= PT \mid PT\backslash \overline{f}$
pure non-dependent types	PS	$::= \circ \mid PS.C \mid P[PS] \mid \&\overline{PS} \mid PS!$
non-dependent types	S	$::= PS \mid PS\backslash \overline{f}$
classes	P	$::= \circ \mid P.C$
values	v	$::= \langle \ell, S \rangle$
access paths	p	$::= v \mid x \mid p.f$
expressions	e	$::= v \mid x \mid e.f \mid x.f = e \mid e_0.m(\overline{e}) \mid e_1;\, e_2$
		$\mid \texttt{new } T \mid (\texttt{view } T)e \mid \texttt{final } T\, x = e_1;\, e_2$
typing environments	Γ	$::= \emptyset \mid \Gamma, x : T \mid \Gamma, p_1 = p_2 \mid \Gamma, Q$

Figure 8: Grammar of the J&$_s$ calculus

not be many masks.

3.3 Translation from the base family to the derived family

Translations in different directions are not always of the same difficulty. Section 3.2 proposes a solution for in-place translation from the derived family (pair) to the base family (base). The complexity of the solution arises mostly because the objects of the Pair class must be translated away. However, as noted in [13], in-place translation in the other direction, that is, from base to pair, should be almost trivial, by simply treating an AST in base as an AST in pair.

To capture the asymmetry, J&$_s$ supports *directional* sharing relationships between types that are possibly masked, represented as $T_1 \rightsquigarrow T_2$, which means that an object of static type T_1 may be applied a view change with target type T_2. In Figure 7, the J&$_s$ compiler may infer the sharing relationship base!.Exp $\rightsquigarrow$ pair!.Exp, (the other direction pair!.Exp $\rightsquigarrow$ base!.Exp does not hold, because of the class Pair), and allows a constant-time in-place translation from base to pair, by a view change from base!.Exp to pair!.Exp.

Note that in this case, the sharing declaration on line 6 in Figure 7 actually induces the following two directional sharing relationships:

$$\texttt{base.Abs!} \rightsquigarrow \texttt{pair.Abs!}$$
$$\texttt{pair.Abs!}\backslash\texttt{e} \rightsquigarrow \texttt{base.Abs!}\backslash\texttt{e}$$

rather than just a bidirectional sharing relationship base.Abs!\e $\leftrightarrow$ pair.Abs!\e.

4 Formal semantics

4.1 Grammar

The grammar of the language is shown in Figure 8. We use the notation $\overline{a}$ for the list $a_1, \ldots, a_n$ for $n \geq 0$. The length of $\overline{a}$ is written $\#(\overline{a})$, and the empty list is written nil. Depending on context, $\overline{a}$ sometimes denotes the set containing all list members. Types with multiple masked fields are written as $T\backslash\overline{f}$.

Every class has a shares clause. If a class C does not share any other class, the type that appears in its shares clause is C itself.

Unlike Featherweight Java [22] or the J& calculus [32], object allocation does not provide any initial values for fields. Every field, in its declaration, specifies an initial value.

Methods have sharing constraints $\overline{Q}$, which are used to type-check view change operations, and to control method overriding. For simplicity, the formal semantics assumes all the classes are known, and the modular type-checking of sharing constraints is not modeled.

In J&$_s$, a value v, that is, a reference, is a pair of a heap location ℓ and a view S, which is a non-dependent exact type that may include masks.

Expressions are mostly standard, with the addition of a view change operation $(\mathtt{view}\,T)e$. Without loss of generality, the receiver of a field assignment is assumed to always be a variable.

The typing environment contains aliasing information about access paths. An entry $p_1 = p_2$ means p_1 and p_2 are aliases, which also have the same run-time view. As in [10], this kind of information is not needed in the static semantics, but just in the soundness proof.

4.2   Auxiliary definitions

Some auxiliary definitions that are straightforward are only informally explained here.

- $\mathsf{super}(P)$ and $\mathsf{share}(P)$ represent the declared supertype and shared type of P. The set of all superclasses of S is $\mathsf{supers}(S)$.

- Well-formedness judgments of types and environments are represented as $\Gamma \vdash T$ and $\vdash \Gamma$ ok.

- $\mathsf{FV}(e)$ and $\mathsf{refs}(e)$ are the sets of free variables and free references in expression e.

- $\mathsf{pure}(T)$ strips all the masks from type T.

4.3   Class lookup

The class table $CT(P)$, defined in Figure 9, returns the declaration for *explicit* class P, or $\bot$ otherwise. In CT-OUT, the premise $Pr = \langle \overline{L}, e \rangle$ indicates that the program consists of a set of top-level class declarations $\overline{L}$ and a "main" expression e, and the "outermost" class $\circ$ contains all the class declarations in $\overline{L}$.

Besides CT, there is also an extended class table CT', which includes implicit classes (but not any top-level intersection types). The class table CT' is needed because sharing constraints must also be checked for implicit classes. In the rule CT$'$-IMP, a declaration of an implicit class $P.C$ is created: the declared supertype T_e is the intersection of the supertypes of all the classes that $P.C$ further binds; the shared type T_s just refers to the class itself with $\mathtt{this.class}.C$. In CT$'$-IMP, $P'' \sqsupseteq_{\mathsf{fb}} P.C$ is a shorthand for $\vdash P.C \sqsubseteq_{\mathsf{fb}} P''$.

4.4   Subclassing

Inheritance among classes is defined in Figure 9. The judgment $\vdash P_1 \sqsubseteq_{\mathsf{sc}} P_2$ states that P_1 is a declared subclass of P_2, and $\vdash P_1.C \sqsubseteq_{\mathsf{fb}} P_2.C$ states that $P_1.C$ further binds $P_2.C$. We write $\vdash P_1 \sqsubseteq P_2$ if P_1 either subclasses or further binds P_2.

The definition uses the mem function, also shown in Figure 9, which returns the set of classes P comprising a pure non-dependent type PS.

4.5   Prefix types

The meaning of a non-dependent prefix type $P[PS]$ is defined using the function $\mathsf{prefix}(P, PS)$:

$$\mathsf{prefix}(P,PS) = \{P' \mid \exists\, C, C' \;.\; \vdash P \sim P' \wedge P.C \in \mathsf{supers}(PS) \wedge P'.C' \in \mathsf{supers}(PS)\}$$

Figure 9: Classes

If $\text{prefixExact}_1(PS) = \text{false}$ (see Section 4.7 for the definition of $\text{prefixExact}_k(T)$), the P-prefix of a pure non-dependent type PS is the intersection of all classes P' where P and P' both inherit a common nested class—that is, P and P' are equivalent under the $\sim$ relation, shown in Figure 9—and PS extends nested classes of both P and P'. This definition ensures that if P is a subtype of P', then $P[PS]$ is equal to $P'[PS]$. If $\text{prefixExact}_1(PS) = \text{true}$, then $P[PS]$ is $(\&\text{prefix}(P,PS))!$, which keeps the exactness of the index.

Note that the index PT of a prefix type $P[PT]$ can only be a pure type, as shown by WF-PRE in Figure 12. Also, as in [32], the J&$_s$ calculus only allow prefix types $P[T]$ where some supertype of T is immediately enclosed by a subclass of P, and more general prefix types can be encoded.

4.6 Member lookup

Figure 9 also shows auxiliary functions for looking up fields and methods: the functions $\mathsf{fields}(P)$ and $\mathsf{methods}(P)$ collect all the field and method declarations from P and its superclasses; $\mathsf{fnames}(\overline{F})$ is the set of all field names in field declarations $\overline{F}$; $\mathsf{ftype}_{decl}(\Gamma,T,f)$ returns the declared type of field f, which might be a type dependent on $\mathtt{this}$; $\mathsf{ftype}(\Gamma,T,f)$ substitute the receiver type for $\mathtt{this.class}$, and returns a reference-only type, if the receiver type has a mask on f; $\mathsf{mtype}(\Gamma,T,m)$ and $\mathsf{mbody}(S,m)$ look up the type and the declaration of a method m.

$\boxed{\Gamma \vdash_{\mathsf{final}} p : T}$

$$\Gamma \vdash_{\mathsf{final}} \langle \ell, S \rangle : S \text{ (F-REF)} \qquad \frac{x : T \in \Gamma}{\Gamma \vdash_{\mathsf{final}} x : T} \text{ (F-VAR)} \qquad \frac{\Gamma \vdash_{\mathsf{final}} p : T \quad T_f = \mathsf{ftype}(\Gamma,T,f)}{\Gamma \vdash_{\mathsf{final}} p.f : T_f} \text{ (F-GET)}$$

$\boxed{\Gamma \vdash p_1 = p_2}$

$$\frac{\begin{array}{c} p_1 = p_2 \in \Gamma \\ p_1 \neq x \quad p_2 \neq x \end{array}}{\Gamma \vdash p_1 = p_2} \text{ (A-ENV)} \qquad \frac{\begin{array}{c} \Gamma \vdash p_1 = p_2 \\ \Gamma \vdash_{\mathsf{final}} p_1.f : T_f \\ \Gamma \vdash_{\mathsf{final}} p_2.f : T_f \end{array}}{\Gamma \vdash p_1.f = p_2.f} \text{ (A-FIELD)} \qquad \frac{\Gamma \vdash_{\mathsf{final}} p : T}{\Gamma \vdash p = p} \text{ (A-REFL)} \qquad \frac{\Gamma \vdash p_2 = p_1}{\Gamma \vdash p_1 = p_2} \text{ (A-SYM)} \qquad \frac{\begin{array}{c} \Gamma \vdash p_1 = p_2 \\ \Gamma \vdash p_2 = p_3 \end{array}}{\Gamma \vdash p_1 = p_3} \text{ (A-TRANS)}$$

$\boxed{\Gamma \vdash T_1 \rightsquigarrow T_2}$

$$\Gamma \vdash T \rightsquigarrow T \text{ (SH-REFL)} \qquad \frac{\Gamma \vdash T_1 \rightsquigarrow T_2 \quad \Gamma \vdash T_2 \rightsquigarrow T_3}{\Gamma \vdash T_1 \rightsquigarrow T_3} \text{ (SH-TRANS)} \qquad \frac{T_1 \rightsquigarrow T_2 \in \Gamma}{\Gamma \vdash T_1 \rightsquigarrow T_2} \text{ (SH-ENV)} \qquad \frac{\Gamma \vdash T_1 \rightsquigarrow T_2}{\Gamma \vdash T_1 \backslash f \rightsquigarrow T_2 \backslash f} \text{ (SH-MASK)}$$

$$\frac{\begin{array}{c} \mathsf{share}(P.C) = P' \backslash \overline{f} \\ \mathsf{fnames}(\mathsf{fields}(P.C) - \mathsf{fields}(P')) = \overline{f'} \end{array}}{\Gamma \vdash P.C! \backslash \overline{f} \backslash \overline{f'} = P'! \backslash \overline{f}} \text{ (SH-DECL)} \qquad \frac{\begin{array}{c} PS_1 = P_1'!.\overline{C_1} \quad PS_2 = P_2'!.\overline{C_2} \\ \forall P_1 . \left(\begin{array}{c} \Gamma \vdash P_1! \leq PS_1 \Rightarrow \\ \exists! P_2 . \exists \overline{f''} \supseteq \overline{f} . \Gamma \vdash P_2! \leq PS_2 \wedge \Gamma \vdash P_1! \backslash \overline{f''} \rightsquigarrow P_2! \backslash \overline{f'} \end{array} \right) \end{array}}{\Gamma \vdash PS_1 \backslash \overline{f} \rightsquigarrow PS_2 \backslash \overline{f'}} \text{ (SH-CLS)}$$

$\boxed{\Gamma \vdash e : T, \Gamma'}$

$$\frac{\Gamma \vdash_{\mathsf{final}} p : T}{\Gamma \vdash p : \mathsf{ptype}(\Gamma,p), \Gamma} \text{ (T-FIN)} \qquad \frac{\Gamma \vdash e : T, \Gamma' \quad \Gamma \vdash T \leq T'}{\Gamma \vdash e : T', \Gamma'} \text{ (T-SUB)} \qquad \frac{\Gamma \vdash e_1 : T_1, \Gamma_1 \quad \Gamma_1 \vdash e_2 : T_2, \Gamma_2}{\Gamma \vdash e_1; e_2 : T_2, \Gamma_2} \text{ (T-SEQ)}$$

$$\Gamma \vdash \mathtt{new}\ T : T!, \Gamma \text{ (T-NEW)} \qquad \frac{\begin{array}{c} \Gamma \vdash e : T, \Gamma' \\ \mathsf{ftype}(\Gamma,T,f) = T_f \end{array}}{\Gamma \vdash e.f : T_f, \Gamma'} \text{ (T-GET)} \qquad \frac{\begin{array}{c} x \notin \mathsf{dom}(\Gamma_1) \quad \Gamma \vdash e_1 : T, \Gamma_1 \\ \Gamma_1, x : T \vdash e_2 : T_2, \Gamma_2 \quad \Gamma_2 = \Gamma_2', x : T' \end{array}}{\Gamma \vdash \mathtt{final}\ T\ x = e_1;\ e_2 : T_2, \Gamma_2'} \text{ (T-LET)}$$

$$\frac{\begin{array}{c} \Gamma \vdash e : T', \Gamma' \\ \Gamma \vdash T' \rightsquigarrow T \end{array}}{\Gamma \vdash (\mathtt{view}\ T)e : T, \Gamma'} \text{ (T-VIEW)} \qquad \frac{\begin{array}{c} \Gamma \vdash e : T_f, \Gamma' \quad \Gamma' \vdash x : T \\ \mathsf{ftype}_{decl}(\Gamma',T',f) = T_f^{decl} \\ T_f^{decl} \{\!\{ \Gamma';\ T/\mathtt{this}! \}\!\} = T_f \end{array}}{\Gamma \vdash x.f = e : T_f, \mathsf{grant}(\Gamma',x.f)} \text{ (T-SET)} \qquad \frac{\begin{array}{c} n = \#(\overline{e}) = \#(\overline{x}) \quad x_0 = \mathtt{this} \quad \Gamma = \Gamma_0 \\ \mathsf{mtype}(\Gamma,T_0^0,m) = (\overline{x} : \overline{T^0}) \rightarrow T_{n+1}^0 \\ \forall i \in 1..n+1, j \in 1..i.\ T_i^{j-1} \{\!\{ \Gamma_{i-1};\ T_{j-1}^{j-1}/x_{j-1}! \}\!\} = T_i^j \\ \forall i \in 0..n.\ \Gamma_i \vdash e_i : T_i^i, \Gamma'_{i+1} \end{array}}{\Gamma \vdash e_0.m(\overline{e}) : T_{n+1}^{n+1}, \Gamma_{n+1}} \text{ (T-CALL)}$$

$\boxed{\Gamma \vdash T \leq T'}$

$$\Gamma \vdash T \leq T \quad \text{(S-REFL)} \qquad \frac{\Gamma \vdash T_1 \leq T_2 \quad \Gamma \vdash T_2 \leq T_3}{\Gamma \vdash T_1 \leq T_3} \text{ (S-TRANS)} \qquad \Gamma \vdash T.C! \leq T!.C \text{ (S-EXACT)}$$

$$\frac{\Gamma \vdash_{\mathsf{final}} p : T \quad \mathsf{pure}(T) = PT}{\Gamma \vdash p.\mathtt{class} \leq PT} \text{ (S-FIN)} \qquad \frac{\Gamma \vdash_{\mathsf{final}} p : T \quad \mathsf{pure}(T) = PT!}{\Gamma \vdash p.\mathtt{class} \approx PT!} \text{ (S-FIN-EXACT)} \qquad \Gamma \vdash T \leq T \backslash f \text{ (S-MASK)}$$

$$\frac{\Gamma \vdash T_1 \leq T_2}{\Gamma \vdash T_1 \backslash f \leq T_2 \backslash f} \text{ (S-SUB-MASK)} \qquad \frac{\Gamma \vdash p_1 = p_2}{\Gamma \vdash p_1.\mathtt{class} \approx p_2.\mathtt{class}} \text{ (S-ALIAS)} \qquad \frac{\Gamma \vdash T \leq P \quad \mathsf{super}(P.C)\{\!\{ \Gamma;\ T/\mathtt{this}! \}\!\} = T'}{\Gamma \vdash T.C \leq T'} \text{ (S-SUP)}$$

$$\frac{\Gamma \vdash T \quad \Gamma \vdash T \unlhd S}{\Gamma \vdash T \leq S} \text{ (S-BOUND)} \qquad \frac{\Gamma \vdash T_1 \leq T_2 \quad \Gamma \vdash T_2.C}{\Gamma \vdash T_1.C \leq T_2.C} \text{ (S-NEST)} \qquad \frac{\Gamma \vdash P[PT.C]}{\Gamma \vdash PT \approx P[PT.C]} \text{ (S-PRE-IN)} \qquad \frac{\Gamma \vdash PT \leq P.C}{\Gamma \vdash PT \leq P[PT].C} \text{ (S-PRE-OUT)}$$

$$\frac{\begin{array}{c} \Gamma \vdash PT_1 \leq PT_2 \\ \Gamma \vdash P[PT_2] \end{array}}{\Gamma \vdash P[PT_1] \leq P[PT_2]} \text{ (S-PRE-1)} \qquad \frac{\begin{array}{c} \vdash P_1 \sim P_2 \vee \vdash P_1 \sqsubseteq_{\mathsf{sc}} P_2 \\ \Gamma \vdash P_1[PT] \quad \Gamma \vdash P_2[PT] \end{array}}{\Gamma \vdash P_1[PT] \approx P_2[PT]} \text{ (S-PRE-2)} \qquad \Gamma \vdash \&\overline{T} \leq T_i \text{ (S-MEET-LB)} \qquad \frac{\forall i.\ \Gamma \vdash T \leq T_i}{\Gamma \vdash T \leq \&\overline{T}} \text{ (S-MEET-G)}$$

Figure 10: Static semantics

15

$$\text{paths}(\circ) = \emptyset$$
$$\text{paths}(T.C) = \text{paths}(T)$$
$$\text{paths}(p.\texttt{class}) = \{p\}$$
$$\text{paths}(P[T]) = \text{paths}(T)$$
$$\text{paths}(\&\overline{T}) = \bigcup_{T_i \in \overline{T}} \text{paths}(T_i)$$
$$\text{paths}(T!) = \text{paths}(T)$$

$$\text{prefixExact}_k(\circ) = \text{false}$$
$$\text{prefixExact}_k(T.C) = \begin{cases} \text{false} & \text{if } k = 0 \\ \text{prefixExact}_{k-1}(T) & \text{otherwise} \end{cases}$$
$$\text{prefixExact}_k(p.\texttt{class}) = \text{true}$$
$$\text{prefixExact}_k(P[T]) = \text{prefixExact}_{k+1}(T)$$
$$\text{prefixExact}_k(\&\overline{T}) = \bigvee_{T_i \in \overline{T}} \text{prefixExact}_k(T_i)$$
$$\text{prefixExact}_k(T!) = \text{true}$$
$$\text{exact}(T) = \text{prefixExact}_0(T)$$

Figure 11: Paths and exactness

$$\frac{CT'(P) \neq \bot}{\Gamma \vdash P} \text{ (WF-SIMP)} \qquad \frac{\Gamma \vdash_{\text{final}} p : T}{\Gamma \vdash p.\texttt{class}} \text{ (WF-FIN)} \qquad \frac{\Gamma \vdash T}{\Gamma \vdash T!} \text{ (WF-EXACT)} \qquad \frac{\Gamma \vdash T \quad T \neq T' \backslash f \quad \Gamma \vdash T \trianglelefteq PS \quad [\texttt{final}] \, T_f \, f \in \text{fields}(PS)}{\Gamma \vdash T \backslash f} \text{ (WF-MASK)}$$

$$\frac{\begin{array}{c} \Gamma \vdash T \\ \Gamma \vdash T \trianglelefteq PS \\ \text{mem}(PS.C) \neq \emptyset \end{array}}{\Gamma \vdash T.C} \text{ (WF-NEST)} \qquad \frac{\begin{array}{c} \Gamma \vdash P \quad \Gamma \vdash T \\ \Gamma \vdash T \trianglelefteq PS \\ \text{pure}(T) = T \\ \text{prefix}(P, PS) \neq \emptyset \end{array}}{\Gamma \vdash P[T]} \text{ (WF-PRE)} \qquad \frac{\begin{array}{c} \forall i. \, \Gamma \vdash T_i \\ \forall p_i, p_j \in \text{paths}(\&\overline{T}). \, \Gamma \vdash p_i = p_j \\ \forall T_i, T_j \in \overline{T}, k \geq 0. \, \text{prefixExact}_k(T_i) \Leftrightarrow \text{prefixExact}_k(T_j) \end{array}}{\Gamma \vdash \&\overline{T}} \text{ (WF-MEET)}$$

Figure 12: Type well-formedness

4.7 Final access paths and exactness

The judgment $\Gamma \vdash_{\text{final}} p : T$, shown in Figure 10, gives the static type bound for final access path p. Note that F-REF does not need a premise, because a reference contains its own type.

The function $\text{paths}(T)$ shown in Figure 11, returns the set of final access paths in the structure of type T.

The function $\text{prefixExact}_k(T)$ shown in Figure 11, is true if the kth prefix of T is an exact type for $k \geq 0$. If $\text{prefixExact}_k(T)$, then $\text{prefixExact}_{k+1}(T)$.

To type-check field accesses in the proof of soundness, the type system keeps track of aliases. The judgment $\Gamma \vdash p_1 = p_2$, defined in Figure 10, states that two final access paths are aliases, and the two references stored in p_1 and p_2 have the same view. Note that A-ENV does *not* allow a judgment of alias with a variable in it, although the environment Γ can have an element in the form of $x = v$ in it.

4.8 Type well-formedness

Type well-formedness is defined in Figure 12. The judgment $\Gamma \vdash T$ states that type T is well-formed in a context Γ. Most of the rules are similar to the J& calculus [32], except WF-EXACT, WF-REF, and WF-MASK, for new kinds of types in J&$_s$.

4.9 Non-dependent bounding types

The judgment $\Gamma \vdash T \trianglelefteq PS$, shown in Figure 13, states that type T has a static type bound PS that is non-dependent and pure, i.e., $\Gamma \vdash \text{pure}(T) \leq PS$. Also, PS is the most specific static type bound of T in the context of Γ. The most interesting rule is BD-FIN, which ensures that the $p_1.\texttt{class}$ and $p_2.\texttt{class}$ have the same type bound, if p_1 and p_2 are aliases in the context of Γ. Therefore a type T might have different type bounds in different typing environments, but the bound is unique in the same environment.

$$\Gamma \vdash PS \trianglelefteq PS \ \text{(BD-SIMP)} \qquad \frac{\Gamma \vdash T \trianglelefteq PS}{\Gamma \vdash T.C \trianglelefteq PS.C} \ \text{(BD-NEST)} \qquad \frac{PS = \&\{PS' \mid \Gamma \vdash p = p' \wedge \Gamma \vdash_{\mathsf{final}} p' : T \wedge \Gamma \vdash T \trianglelefteq PS'\}}{\Gamma \vdash p.\mathsf{class} \trianglelefteq PS} \ \text{(BD-FIN)}$$

$$\frac{\Gamma \vdash T \trianglelefteq PS}{\Gamma \vdash P[T] \trianglelefteq P[PS]} \ \text{(BD-PRE)} \qquad \frac{\forall i.\ \Gamma \vdash T_i \trianglelefteq PS_i}{\Gamma \vdash \&\overline{T} \trianglelefteq \&\overline{PS}} \ \text{(BD-MEET)} \qquad \frac{\Gamma \vdash T \trianglelefteq PS \quad T \neq PS'}{\Gamma \vdash T! \trianglelefteq PS} \ \text{(BD-EXACT)} \qquad \frac{\Gamma \vdash \mathsf{pure}(T) \trianglelefteq PS}{\Gamma \vdash T \trianglelefteq PS} \ \text{(BD-PURE)}$$

Figure 13: Type bounds

$$\circ \{\!\{\Gamma;\ T_x/x\}\!\} = \circ$$

$$T.C\{\!\{\Gamma;\ T_x/x\}\!\} = T\{\!\{\Gamma;\ T_x/x\}\!\}.C$$

$$v.\mathsf{class}\{\!\{\Gamma;\ T_x/x\}\!\} = v.\mathsf{class}$$

$$\frac{x \neq y}{y.\mathsf{class}\{\!\{\Gamma;\ T_x/x\}\!\} = y.\mathsf{class}}$$

$$x.\mathsf{class}\{\!\{\Gamma;\ T_x/x\}\!\} = \mathsf{pure}(T_x)$$

$$\frac{p.\mathsf{class}\{\!\{\Gamma;\ T_x/x\}\!\} = p'.\mathsf{class}}{p.f.\mathsf{class}\{\!\{\Gamma;\ T_x/x\}\!\} = p'.f.\mathsf{class}}$$

$$\frac{\begin{array}{c} p.\mathsf{class}\{\!\{\Gamma;\ T_x/x\}\!\} = T_p \\ T_p \neq p'.\mathsf{class} \\ \mathsf{ftype}(\Gamma, T_p, f) = T_f \end{array}}{p.f.\mathsf{class}\{\!\{\Gamma;\ T_x/x\}\!\} = \mathsf{pure}(T_f)}$$

$$\frac{T\{\!\{\Gamma;\ T_x/x\}\!\} = T'}{P[T]\{\!\{\Gamma;\ T_x/x\}\!\} = P[T']}$$

$$\frac{\forall i.\ T_i\{\!\{\Gamma;\ T_x/x\}\!\} = T_i'}{\&\overline{T}\{\!\{\Gamma;\ T_x/x\}\!\} = \&\overline{T'}}$$

$$\frac{T\{\!\{\Gamma;\ T_x/x\}\!\} = T'}{T!\{\!\{\Gamma;\ T_x/x\}\!\} = T'!}$$

$$\frac{T\{\!\{\Gamma;\ T_x/x\}\!\} = T'}{T\backslash f\{\!\{\Gamma;\ T_x/x\}\!\} = T'\backslash f}$$

Figure 14: Type substitution

4.10 Type substitution

The rules for type substitution are shown in Figure 14. Type substitution $T\{\!\{\Gamma;\ T_x/x\}\!\}$ substitutes $\mathsf{pure}(T_x)$ for $x.\mathsf{class}$ in T in the context of Γ. The typing context Γ is used to look up field types when substituting a non-dependent class into a field-path dependent class. T_x should be well-formed and a subtype of x's declared type.

For type safety, type substitutions on the right-hand side of a field assignment or on the parameters of method calls must preserve the exactness of the declared type. Therefore, only values from the family that is compatible with the receiver are assigned to fields or passed to method code. Exactness-preserving type substitution $T\{\!\{\Gamma;\ T_x/x!\}\!\}$ is defined as follows:

$$\frac{\begin{array}{c} T\{\!\{\Gamma;\ T_x/x\}\!\} = T' \\ \forall k.\ \mathsf{prefixExact}_k(T) \Rightarrow \mathsf{prefixExact}_k(T') \end{array}}{T\{\!\{\Gamma;\ T_x/x!\}\!\} = T'}$$

4.11 Sharing relationships

The directional sharing judgment $\Gamma \vdash T_1 \rightsquigarrow T_2$, shown in Figure 10, states that a value of type T_1 can be transformed to T_2 through a view change, that is, the sharing relationship $T_1 \rightsquigarrow T_2$ is valid. A bidirectional sharing judgment $\Gamma \vdash T_1 \leftrightarrow T_2$ is sugar for a pair of directional sharing judgments. SH-DECL collects sharing relationships from class declarations, and for simplicity, it does not model the inference of asymmetric sharing relationships, described in Section 3.3. SH-CLS states that for two types to be shared, every subclass of the source type must share a unique subclass of the target type, with appropriate masks. The uniqueness is necessary for the view function to set the type component in the generated reference. Moreover, the two types in SH-CLS have to be both nested in simple exact types, and therefore we can enumerate all their subclasses in the locally closed worlds.

4.12 Typing rules

Expression typing rules are shown in Figure 10. Evaluation of an expression can update type bindings of some variables. For example, assignment to a masked field removes the mask. Therefore, typing judgments

are of the form $\Gamma \vdash e : T, \Gamma'$, stating that e has type T in context Γ, and that after evaluating e, an updated environment Γ' exists instead. Typing judgments are sometimes written as $\Gamma \vdash e : T$, if the updated typing environment is unused.

Masks can be introduced by T-SUB and S-MASK, and they can be removed by field assignments. The following auxiliary function grant removes the mask, that is, granting the access to a field: $\mathsf{grant}(\Gamma, e)$ generates an environment from Γ, where accesses to $x.f$ are enabled, for $e = x.f$.

$$\mathsf{grant}(\Gamma, e) = \begin{cases} \Gamma_1, x : T & \text{if } \Gamma = \Gamma_1, x : T \backslash f \wedge e = x.f \\ \Gamma & \text{otherwise} \end{cases}$$

The typing of a final access path p is complicated by the fact that $p.\mathtt{class}$ is pure, and therefore p might not have $p.\mathtt{class}$ as its type. The auxiliary function ptype, shown below, gives a proper type to p, and it also avoids dependent classes that have paths starting with a variable, if the typing environment contains the information about the variable on the stack.

$$\mathsf{ptype}(\Gamma, p) = \begin{cases} \mathsf{ptype}(\Gamma, v.\overline{f}) & \text{if } p = x.\overline{f} \wedge x = v \in \Gamma \\ p.\mathtt{class} \backslash \overline{f} & \text{if } \Gamma \vdash_{\mathsf{final}} p : PT \backslash \overline{f} \end{cases}$$

In addition, A-ENV does not give aliasing information that includes variables. Therefore, in the proof of soundness, types that depend on paths starting with variables can be avoided all together, because the stack σ contains the value for every variable.

4.13 Subtyping

Subtyping rules are defined in Figure 10. The judgment $\Gamma \vdash T \leq T'$ states that T is a subtype of T' in context Γ, and type equivalence $\Gamma \vdash T \approx T'$ is sugar for a pair of subtyping judgments.

The J&$_s$ type system clearly distinguishes sharing relationships between types from subtyping relationships. Subtyping rules in J&$_s$ do not depend on any sharing judgment. In particular, $\Gamma \vdash T_1 \rightsquigarrow T_2$ does not imply $\Gamma \vdash T_1 \leq T_2$.

4.14 Program typing

Program typing rules are shown in Figure 15. P-OK is the rule for a program to be well-formed; L-OK, F-OK, and M-OK are the rules for a class, a field, and a method to be well-formed, respectively.

Q-OK gives the rule for a sharing constraint to be well-formed: the two types, interpreted in the context of the current containing class, should be shared. Also, whenever a sharing constraint in an inherited method no longer holds, the method should be overridden.

4.15 Operational semantics

A small-step semantics for J&$_s$ is shown in Figures 16 and 17. A stack σ is a function mapping variable names x to values v. A heap H is a function mapping tuples $\langle \ell, P, f \rangle$ of memory locations, classes, and field names to values v. The J&$_s$ calculus follows the approach of J\mask [36], where uninitialized fields do not exist in the heap. A J&$_s$ object might have multiple versions of a field f, when the interpreted types of f are not shared for different views of the object. The class P in the domain of the heap to indicate which copy of the possibly unshared field f is used. Given a view $P'.C! \backslash \overline{f'}$ of ℓ and a field f, the indicator class P is

$$\frac{\circ \vdash \overline{L} \text{ ok} \quad \emptyset \vdash e:T \quad \emptyset \vdash T \quad \sqsubseteq^+ \text{ acyclic}}{\vdash \langle \overline{L}, e \rangle \text{ ok}} \tag{P-OK}$$

$$\frac{S_1 = T_1\{\!\{\emptyset;\, P!/\texttt{this}!\}\!\} \quad S_2 = T_2\{\!\{\emptyset;\, P!/\texttt{this}!\}\!\} \quad \emptyset \vdash S_1 \rightsquigarrow S_2}{P \vdash T_1 \rightsquigarrow T_2 \text{ ok}} \tag{Q-OK}$$

$$\frac{\begin{array}{c} \forall C'.\ CT'(P.C.C') \neq \bot \Rightarrow P.C \vdash CT'(P.C.C') \text{ ok} \\ P.C \vdash \overline{F} \text{ ok} \quad P.C \vdash \overline{M} \text{ ok} \quad P \vdash T \text{ super ok} \\ \forall P_i \in \text{supers}(P.C) \backslash \{P.C\}.\ \vdash P.C \text{ conforms to } P_i \\ T' = P' \backslash \overline{f} \quad \vdash P.C \sqsubseteq^*_{\text{fb}} P' \\ \forall f.\ \texttt{final } T_f\, f\, \ldots \in \text{fields}(P') \Rightarrow f \notin \overline{f} \\ \forall f.\ \bigl(\ldots\, T_f\, f\, \ldots \in \text{fields}(P') \wedge f \notin \overline{f} \Rightarrow\, \vdash T_f\{\!\{\emptyset;\, P'/\texttt{this}\}\!\} \leftrightarrow T_f\{\!\{\emptyset;\, P.C/\texttt{this}\}\!\} \bigr) \end{array}}{P \vdash \texttt{class } C \texttt{ extends } T \texttt{ shares } T'\ \{\overline{L}\,\overline{F}\,\overline{M}\} \text{ ok}} \tag{L-OK}$$

$$\frac{T \neq \circ \quad \texttt{this}:P \vdash T \quad \text{paths}(T) \subseteq \{\texttt{this}\} \quad \neg\text{exact}(T)}{P \vdash T \text{ super ok}}$$

$$\frac{\begin{array}{c} CT'(P) = \texttt{class } C \texttt{ extends } T_e \texttt{ shares } T_s\ \{\overline{L}\,\overline{F}\,\overline{M}\} \\ CT'(P') = \texttt{class } C' \texttt{ extends } T_e' \texttt{ shares } T_s'\ \{\overline{L'}\,\overline{F'}\,\overline{M'}\} \\ \forall i,j.\ \Bigl(L_i = \texttt{class } D \texttt{ extends } T_i\ \ldots\ \{\ldots\} \wedge L_j' = \texttt{class } D \texttt{ extends } T_j'\ \ldots\ \{\ldots\} \Bigr) \Rightarrow \texttt{this}:P \vdash T_i \leq T_j' \\ (\text{fnames}(\overline{F}) \cap \text{fnames}(\overline{F'})) = \emptyset \\ \forall i,j.\ \Bigl(M_i = T_{n+1}\, m(\overline{T}\,\overline{x})\ \{e\} \wedge M_j' = T_{n+1}'\, m(\overline{T'}\,\overline{x'})\ \{e'\} \Bigr) \Rightarrow P \vdash M_i \text{ overrides } M_j' \\ \forall i.\ M_i' = \ldots \texttt{ sharing } \overline{Q}\ \ldots \wedge P \not\vdash \overline{Q} \text{ ok} \Rightarrow \exists j.\ P \vdash M_j \text{ overrides } M_i' \end{array}}{\vdash P \text{ conforms to } P'}$$

$$\frac{\begin{array}{c} M = T_{n+1}\, m(\overline{T}\,\overline{x})\ \ldots\ \{e\} \\ M' = T_{n+1}'\, m(\overline{T'}\,\overline{x'})\ \ldots\ \{e'\} \\ \#(\overline{x}) = \#(\overline{x'}) = \#(\overline{y}) \quad \overline{y} \cap (\overline{x} \cup \overline{x'}) = \emptyset \\ \Gamma = \texttt{this}:P, \overline{y}:\overline{T}\{\overline{y}/\overline{x}\} \quad \vdash \Gamma \text{ ok} \\ \Gamma \vdash \overline{T}\{\overline{y}/\overline{x}\} \approx \overline{T'}\{\overline{y}/\overline{x'}\} \quad \Gamma \vdash T_{n+1}\{\overline{y}/\overline{x}\} \approx T_{n+1}'\{\overline{y}/\overline{x'}\} \end{array}}{P \vdash M \text{ overrides } M'}$$

$$\frac{\begin{array}{c} \text{fnames}(\text{fields}(P)) = \overline{f} \quad \Gamma = \texttt{this}:P\backslash \overline{f} \\ \Gamma \vdash T \quad \text{paths}(T) \subseteq \{\texttt{this}\} \quad \text{exact}(T) = \texttt{false} \quad \Gamma \vdash e:T,\Gamma \end{array}}{P \vdash [\texttt{final}]\ T\ f = e \text{ ok}} \tag{F-OK}$$

$$\frac{\begin{array}{c} P \vdash \overline{Q} \text{ ok} \quad \Gamma = \texttt{this}:P, \overline{x}:\overline{T} \\ \vdash \Gamma \text{ ok} \quad \Gamma \vdash T_{n+1} \quad \Gamma \vdash e:T_{n+1}, \Gamma_r \quad \text{FV}(e) \subseteq \{x_0, \overline{x}\} \\ n = \#(\overline{x}) \quad x_0 = \texttt{this} \\ \forall i \in 1..n+1.\ \text{paths}(T_i) \in \{x_0, \ldots, x_{i-1}\} \end{array}}{P \vdash T_{n+1}\, m(\overline{T}\,\overline{x}) \texttt{ sharing } \overline{Q}\ \{e\} \text{ ok}} \tag{M-OK}$$

Figure 15: Program typing

obtained with the following auxiliary function fclass:

$$\text{fclass}(P'.C, f) = \begin{cases} P'.C & \text{if } P'.C = \text{share}(P'.C) \\ P'.C & \text{if } f \in \text{fnames}(\text{fields}(P'.C) - \text{fields}(\text{share}(P'.C))) \\ & \quad \vee \text{share}(P'.C) = T\backslash f \\ \text{fclass}(\text{pure}(\text{share}(P'.C)), f) & \text{otherwise} \end{cases}$$

Stack updates and heap updates are represented as $\sigma[x := v]$ and $H[\langle \ell, P, f \rangle := v]$ respectively. We do not explicitly model popping off stack frames or garbage collection on the heap.

A reference set R, which contains all the references v that have been generated during evaluation, no matter whether they are reachable from the stack or not, is also part of the evaluation configuration e, σ, H, R. The set R is only for the proof of soundness: it prevents us from losing path equalities needed in the proof.

The evaluation rules (Figure 17) take the form $e, \sigma, H, R \longrightarrow e', \sigma', H', R'$. We overload the function grant to work on σ: it updates the type components of values $\langle \ell, S \rangle$ stored in σ with appropriate annotations.

R-SET assigns a value to a field, and may remove the mask on field $x.f$. We overload the function grant,

$$
\begin{array}{llll}
\text{stacks} & \sigma & ::= \emptyset \mid \sigma, x \mapsto v \\
\text{heaps} & H & ::= \emptyset \mid H, \langle \ell, \overline{P}, f \rangle \mapsto v \\
\text{reference sets} & R & ::= \emptyset \mid R, v \\
\text{evaluation contexts} & E & ::= [\cdot] \mid E.f \mid x.f = E \\
& & \mid x \mid E;\ e \mid (\texttt{view}\ TE)e \mid (\texttt{view}\ S)E \\
& & \mid E.m(\overline{e}) \mid v.m(\overline{v}, E, \overline{e}) \\
& & \mid \texttt{new}\ TE \mid \texttt{final}\ TE\ x = e_1;\ e_2 \mid \texttt{final}\ S\ x = E;\ e_2 \\
\text{type evaluation contexts} & TE & ::= TE.C \mid E.\texttt{class} \mid P[TE] \\
& & \mid \&(\overline{S}, TE, \overline{T}) \mid TE! \\
& & \mid TE \backslash \overline{f} \mid v.\texttt{class}
\end{array}
$$

Figure 16: Definitions for operational semantics

defined in Section 4.12, to work on σ: the type components of values $\langle \ell, S \rangle$ stored in σ are updated with appropriate annotations.

R-ALLOC is the rule for `new` expressions, which initialize all the fields according to field initializers collected from class declarations. Note that other views' copies of unshared fields are not initialized right after the object is created. Masked types ensure that after a view change, those possibly uninitialized fields are not accessed.

The order of evaluation is captured by an evaluation context E. There is also a type evaluation context TE for evaluating dependent types. A fully evaluated type will always be non-dependent, since $\langle \ell, S \rangle.\texttt{class}$ can evaluate to S.

Since values can be viewed with different type S, the function view takes a value v and a non-dependent type S, and returns a different value consisting of the same location, with its view changed to be compatible with S. This function is used for field accesses and view changing operations, and the rule SH-CLS ensures that there is always a single view to change to. Note that the view in a reference is always an exact type.

$$
\text{view}(\langle \ell, P'! \backslash \overline{f'} \rangle, PS \backslash \overline{f}) =
\begin{cases}
\langle \ell, P'! \backslash \overline{f} \rangle & \vdash P'! \backslash \overline{f'} \leq PS \backslash \overline{f} \\
\langle \ell, P! \backslash \overline{f} \rangle & \exists! P\ .\ \exists \overline{f''} \supseteq \overline{f'}\ .\ \vdash P! \leq PS \wedge \vdash P'! \backslash \overline{f''} \rightsquigarrow P! \backslash \overline{f}
\end{cases}
$$

In R-ALLOC, the judgment $H \vdash \ell\ \texttt{fresh}$ means that there is no tuple $\langle \ell, P, f \rangle$ in the domain of H, for any P and any f.

5 Soundness

5.1 Runtime typing environments

In the proof of soundness, run-time values are typed using a typing environment $\lfloor \sigma, H, R \rfloor$ constructed from the stack, the heap, and the reference set. Although references of different shared types might address the same heap location, the construction of $\lfloor \sigma, H, R \rfloor$ ensures that they are *not* included as aliases, which is necessary for S-ALIAS to hold. See Figure 18 for details.

5.2 Well-formed configurations

The well-formedness of a configuration e, σ, H, R is given in Figure 19.

A configuration e, σ, H, R is well-formed, if every value pointed to by an unmasked field has (or can be transformed to) a view compatible with the type of its container. This ensures that field accesses always succeed. Also, all the free variables in e must be in the domain of σ, and all the references in e must be in R.

$$\boxed{e,\sigma,H,R \longrightarrow e',\sigma',H',R'}$$

$$\frac{e,\sigma,H,R \longrightarrow e',\sigma',H',R'}{E[e],\sigma,H,R \longrightarrow E[e'],\sigma',H',R'} \qquad\qquad \text{(R-CONG)}$$

$$\frac{\sigma(x) = v}{x,\sigma,H,R \longrightarrow v,\sigma,H,R} \qquad\qquad \text{(R-VAR)}$$

$$\frac{y \notin \mathsf{dom}(\sigma) \quad \sigma' = \sigma[y := v_1]}{\mathtt{final}\ S\ x = v_1;\ e_2,\sigma,H,R \longrightarrow e_2\{y/x\},\sigma',H,R} \qquad\qquad \text{(R-LET)}$$

$$\frac{H(\ell,\mathsf{fclass}(P,f),f) = v \quad S = \mathsf{ftype}(\emptyset,P!\backslash\overline{f'},f) \quad R' = R,\mathsf{view}(v,S)}{\langle \ell,P!\backslash\overline{f'}\rangle.f,\sigma,H,R \longrightarrow \mathsf{view}(v,S),\sigma,H,R'} \qquad\qquad \text{(R-GET)}$$

$$\frac{\sigma(x) = \langle \ell,P!\backslash\overline{f'}\rangle \quad \sigma' = \mathsf{grant}(\sigma,x.f) \quad H' = H[\langle \ell,\mathsf{fclass}(P,f),f\rangle := v] \quad R' = R,\langle \ell,P!\backslash(\overline{f'}-f)\rangle}{x.f = v,\sigma,H,R \longrightarrow v,\sigma',H',R'} \qquad\qquad \text{(R-SET)}$$

$$\frac{\mathsf{mbody}(S,m) = T_{n+1}\ m(\overline{T}\ \overline{x})\ \{e\} \quad n = \#(\overline{v}) = \#(\overline{x}) \quad y_0,y_1,\dots,y_n \notin \mathsf{dom}(\sigma) \quad \sigma' = \sigma[y_0 := \langle \ell,S\rangle,\overline{y} := \overline{v}]}{\langle \ell,S\rangle.m(\overline{v}),\sigma,H,R \longrightarrow e\{y_0/\mathtt{this},\overline{y}/\overline{x}\},\sigma',H,R} \qquad\qquad \text{(R-CALL)}$$

$$\frac{H \vdash \ell\ \mathtt{fresh} \quad x \notin \mathsf{dom}(\sigma) \quad \mathsf{fields}(S) = [\mathtt{final}]\ \overline{T_f}\ \overline{f} = \overline{e} \quad v = \langle \ell,S!\backslash\overline{f}\rangle \quad R' = R,v}{\mathtt{new}\ S,\sigma,H,R \longrightarrow \mathtt{final}\ S!\backslash\overline{f}\ x = v;\ x.\overline{f} = \overline{e}\{x/\mathtt{this}\};\ x,\sigma,H,R'} \qquad\qquad \text{(R-ALLOC)}$$

$$v;\ e,\sigma,H,R \longrightarrow e,\sigma,H,R \qquad\qquad \text{(R-SEQ)}$$

$$\frac{R' = R,\mathsf{view}(v,S)}{(\mathtt{view}\ S)v,\sigma,H,R \longrightarrow \mathsf{view}(v,S),\sigma,H,R'} \qquad\qquad \text{(R-VIEW)}$$

Figure 17: Small-step operational semantics

5.3 Soundness

We prove soundness using subject reduction and progress [51]. The proof of subject reduction is the hard part. There are two notable issues:

- The typing environment $\lfloor \sigma',H',R' \rfloor$ after an evaluation step is *not* a simple superset of the environment $\lfloor \sigma,H,R \rfloor$ before the step, because some field assignment might update variable typing (R-SET). However the updated type is always a subtype of the original type. Therefore, we define the environment extension as follows: Γ_2 extends Γ_1 if

  - For every $p_1 = p_2 \in \Gamma_1$, there is $p_1 = p_2 \in \Gamma_2$;
  - For every $T_1 \rightsquigarrow T_2 \in \Gamma_1$, there is $T_1 \rightsquigarrow T_2 \in \Gamma_2$;
  - For every $x : T \in \Gamma_1$, there is $x : T' \in \Gamma_2$ and $\Gamma_2 \vdash T' \leq T$.

- Method calls are evaluated by adding a set of fresh variable bindings (essentially, a new stack frame) into σ, rather than directly substituting actual parameters for formal arguments.

We first prove some lemmas about extensions of typing environments:

Lemma 5.1 *If $e,\sigma,H,R \longrightarrow e',\sigma',H',R'$, then $\lfloor \sigma',H',R' \rfloor$ is an extension of $\lfloor \sigma,H,R \rfloor$.*

PROOF: By induction on the operational semantic derivation.

21

$$\frac{\sigma(x) = \langle \ell, S \rangle}{x{:}S, x = \langle \ell, S \rangle \in \lfloor \sigma, H, R \rfloor}$$

$$\frac{\langle \ell, PS \backslash \overline{f} \rangle \in R \quad \langle \ell, PS \backslash \overline{f'} \rangle \in R}{\langle \ell, PS \backslash \overline{f} \rangle = \langle \ell, PS \backslash \overline{f'} \rangle \in \lfloor \sigma, H, R \rfloor}$$

$$\frac{\begin{array}{c} \langle \ell, S \rangle \in R \quad S = P! \backslash \overline{f'} \\ v = \mathsf{view}(H(\ell, \mathsf{fclass}(P, f), f), T_f) \\ \mathsf{ftype}(\emptyset, S, f) = T_f \quad \mathtt{final}\ T_f^{decl}\ f = e \in \mathsf{fields}(S) \end{array}}{\langle \ell, S \rangle.f = v \in \lfloor \sigma, H, R \rfloor}$$

Figure 18: Runtime typing environments

$$\frac{\begin{array}{c} \mathsf{FV}(e) \subseteq \mathsf{dom}(\sigma) \quad \mathsf{refs}(e) \subseteq R \\ \forall \langle \ell, P! \backslash \overline{f'} \rangle \in R.\ \forall f \notin \overline{f'}.\ \exists \ell'\ .\ \exists S'\ .\ \left(\begin{array}{c} H(\ell, \mathsf{fclass}(P, f), f) = \langle \ell', S' \rangle \wedge \\ (\vdash S' \leq \mathsf{ftype}(\emptyset, S, f) \vee \exists S''\ .\ \vdash S'' \leq \mathsf{ftype}(\emptyset, S, f) \wedge \vdash S' \rightsquigarrow S'') \end{array} \right) \end{array}}{\vdash e, \sigma, H, R} \quad \text{(CONFIG)}$$

Figure 19: Well-formed configurations

- R-CONG

  By the induction hypothesis.

- R-VAR, and R-SEQ

  Trivial, because $\sigma' = \sigma$, $H' = H$, and $R' = R$ in these cases.

- R-LET

  Then $H' = H$, $R' = R$, and $\sigma' = \sigma[y := v_1]$ where y is fresh. We will have $\lfloor \sigma', H, R \rfloor = \lfloor \sigma, H, R \rfloor, y{:}S, y = v_1$, and $\lfloor \sigma', H, R \rfloor$ is an extension of $\lfloor \sigma, H, R \rfloor$.

- R-GET

  Then $\sigma' = \sigma$, $H' = H$, and R' might be a superset of R. By the definition of run-time typing environments, $\lfloor \sigma, H, R' \rfloor$ can only contain some aliasing information other than elements in $\lfloor \sigma, H, R \rfloor$. Thus $\lfloor \sigma, H, R' \rfloor$ is an extension of $\lfloor \sigma, H, R \rfloor$.

- R-SET

  Then $\sigma' = \mathsf{grant}(\sigma, x.f)$, $S = P! \backslash \overline{f'}$, $H' = H[\langle \ell, \mathsf{fclass}(P, f), f \rangle := v]$, and R' might be a superset of R. Compared to $\lfloor \sigma, H, R \rfloor$, the new environment $\lfloor \sigma', H', R' \rfloor$ adds some aliasing relationships, and it might update the type binding of x from $T \backslash f$ to T for some type T. By S-MASK, $\lfloor \sigma', H', R' \rfloor \vdash T \leq T \backslash f$. Thus $\lfloor \sigma', H', R' \rfloor$ is an extension of $\lfloor \sigma, H, R \rfloor$.

- R-CALL

  Similar to the case of R-LET.

- R-ALLOC

  Similar to the case of R-GET.

- R-VIEW

  Similar to the case of R-GET.

$\square$

Lemma 5.2 *If Γ_2 is an extension of Γ_1, then all of the following hold:*

- *If $\Gamma_1 \vdash T$, then $\Gamma_2 \vdash T$.*

- *If $\Gamma_1 \vdash T \trianglelefteq S$, then $\Gamma_2 \vdash T \trianglelefteq S$.*

- *If $\mathsf{ftype}(\Gamma_1, T, f) = T_f$, then $\mathsf{ftype}(\Gamma_2, T, f) = T_f$.*

- *If $\mathsf{mtype}(\Gamma_1, T, m) = (\bar{x}:\overline{T}) \to T_{n+1}$, then $\mathsf{mtype}(\Gamma_2, T, m) = (\bar{x}:\overline{T}) \to T_{n+1}$.*

- *If $\Gamma_1 \vdash p = p'$, then $\Gamma_2 \vdash p = p'$.*

- *if $\mathsf{ptype}(\Gamma_1, p) = T$, then $\mathsf{ptype}(\Gamma_2, p) = T$.*

- *If $T\{\!\{\Gamma_1;\ T_x/x\}\!\} = T'$, then $T\{\!\{\Gamma_2;\ T_x/x\}\!\} = T'$.*

- *If $\Gamma_1 \vdash T_1 \leq T_2$, then $\Gamma_2 \vdash T_1 \leq T_2$.*

- *If $\Gamma_1 \vdash v : T$, then $\Gamma_2 \vdash v : T$.*

PROOF: The proof is by induction on the derivation of the appropriate judgment. $\square$

Note that although Lemma 5.2 states that value typing does not change during the execution, it does not directly apply to more general expression typing, because variables can change their types during the execution.

The following lemma is about type substitution of dependent classes.

Lemma 5.3 *If $x:T_x \in \Gamma$, $\mathsf{paths}(T_x) = \emptyset$, $\Gamma \vdash S_x! \leq T_x$, and $\Gamma \vdash e : T$, then $\Gamma\{\!\{\emptyset;\ S_x!/x\}\!\} \vdash e : T\{\!\{\emptyset;\ S_x!/x\}\!\}$.*

PROOF: The proof is by induction on the typing derivation $\Gamma \vdash e : T$. $\square$

Type substitution is generalized to the typing environment Γ, where substitution is applied to the type of every variable in Γ.

The following lemma establishes connections between type substitutions and subtyping relationships. It is generally used together with Lemma 5.3.

Lemma 5.4 *If $\Gamma \vdash T_1 \leq T_2$, then $\Gamma \vdash T\{\!\{\Gamma;\ T_1/x\}\!\} \leq T\{\!\{\Gamma;\ T_2/x\}\!\}$.*

PROOF: By induction on the structure of T. Note that there is no first-order function type, and therefore a type substitution does not occur at a contravariant position. $\square$

Lemma 5.5 *If $\vdash p, \sigma, H, R$, and $\lfloor \sigma, H, R \rfloor \vdash_{\mathsf{final}} p : T$, and $p, \sigma, H, R \longrightarrow p', \sigma', H', R'$, and $\lfloor \sigma', H', R' \rfloor \vdash_{\mathsf{final}} p' : T'$, then $\lfloor \sigma, H, R \rfloor \vdash \mathsf{ptype}(\lfloor \sigma, H, R \rfloor, p) \approx \mathsf{ptype}(\lfloor \sigma, H, R \rfloor, p')$.*

PROOF: There are the following three cases:

- $p = x.\bar{f}$, where $\bar{f}$ may be empty

  Then by F-GET and F-VAR, $x:T_x \in \lfloor \sigma, H, R \rfloor$, which implies $\sigma(x) = \langle \ell, T_x \rangle$ and $x = \langle \ell, T_x \rangle \in \lfloor \sigma, H, R \rfloor$ by the definition of $\lfloor \sigma, H, R \rfloor$, and by R-VAR and R-CONG, $p' = \langle \ell, T_x \rangle.\bar{f}$. By the definition of ptype, we have $\mathsf{ptype}(\lfloor \sigma, H, R \rfloor, p) = \mathsf{ptype}(\lfloor \sigma, H, R \rfloor, p')$. Therefore, by S-REFL, $\lfloor \sigma, H, R \rfloor \vdash \mathsf{ptype}(\lfloor \sigma, H, R \rfloor, p) \approx \mathsf{ptype}(\lfloor \sigma, H, R \rfloor, p')$.

- $p = v$

  Vacuously true, since it cannot make any progress.

- $p = v.f.\overline{f'}$, where $v = \langle \ell, P! \backslash \overline{f''} \rangle$

  Then by R-GET and R-CONG, we have $p' = v'.\overline{f'}$, $H(\ell, \mathsf{fclass}(P, f), f) = v''$, and $v' = \mathsf{view}(v'', \mathsf{ftype}(\emptyset, P! \backslash \overline{f''}, f))$. Therefore, by the definition of $\lfloor \sigma, H, R \rfloor$, we have $v.f = v' \in \lfloor \sigma, H, R \rfloor$. By A-ENV, $\lfloor \sigma, H, R \rfloor \vdash v.f = v'$. By A-FIELD, $\lfloor \sigma, H, R \rfloor \vdash p = p'$. By S-ALIAS, $\lfloor \sigma, H, R \rfloor \vdash p.\mathtt{class} \approx p'.\mathtt{class}$. Now we need to show that T and T' have the same set of masks. There are again two cases:

  - $\overline{f'} = \emptyset$

    By the definition of the auxiliary function view, T' has the same set of masks as $\mathsf{ftype}(\emptyset, P! \backslash \overline{f''}, f)$, which is T.

  - $\overline{f'} \neq \emptyset$

    Then T and T' have the same set of masks, because they are the type of the same field (the last field in $\overline{f'}$).

$\square$

In the proof of subject reduction, Lemma 5.2 proves the case where T-SUB is the last derivation; Lemma 5.3 is used to prove type preservation for the evaluation of a method call.

Lemma 5.6 *(Subject reduction)* *If* $\vdash e, \sigma, H, R$, *and* $\lfloor \sigma, H, R \rfloor \vdash e : T$, *and* $e, \sigma, H, R \longrightarrow e', \sigma', H', R'$, *then* $\vdash e', \sigma', H', R'$, *and* $\lfloor \sigma', H', R' \rfloor \vdash e' : T$.

PROOF: The proof is by induction on the typing derivation $\lfloor \sigma, H, R \rfloor \vdash e : T, \Gamma'$.

In order to handle the context-sensitivity of the type system, the induction hypothesis is strengthened to:

If $\vdash e, \sigma, H, R$, and $\lfloor \sigma, H, R \rfloor \vdash e : T, \Gamma$, and $e, \sigma, H, R \longrightarrow e', \sigma', H', R'$, then $\vdash e', \sigma', H', R'$, and $\lfloor \sigma', H', R' \rfloor \vdash e' : T, \Gamma'$, and Γ' is an extension of Γ

In order to prove $\vdash e', \sigma', H', R'$, we only need to consider evaluation steps where $H' \neq H$ or $R' \neq R$. The congruence rule R-CONG can be proved by the induction hypothesis directly. The remaining cases are R-GET, R-SET, R-ALLOC, and R-VIEW, and the interesting cases are R-GET, R-VIEW, and R-SET, where R is updated with a value as the result of view or a new value with one of its masks removed. The proof is according to L-OK: when two classes are declared as being shared, their unmasked fields have shared types.

Now it remains to prove $\lfloor \sigma', H', R' \rfloor \vdash e' : T, \Gamma'$, where Γ' is an extension of Γ.

First, if the last derivation uses T-SUB, then it follows from the induction hypothesis and Lemma 5.2. Now it is only necessary to consider derivations that do not end with T-SUB. Consider different cases of e:

- $e = v$.

  Vacuously true, since evaluation cannot continue.

- $e = x$

  The evaluation is $x, \sigma, H, R \longrightarrow v, \sigma, H, R$ by R-VAR, where $\sigma(x) = v$. Suppose $v = \langle \ell, P! \backslash \overline{f} \rangle$, and then by T-FIN and the definition of $\lfloor \sigma, H, R \rfloor$, we have $T = v.\mathtt{class} \backslash \overline{f}$ and $\lfloor \sigma, H, R \rfloor \vdash x : T, \lfloor \sigma, H, R \rfloor$. Then by F-REF and T-FIN, $\lfloor \sigma, H, R \rfloor \vdash v : v.\mathtt{class} \backslash \overline{f}, \lfloor \sigma, H, R \rfloor$.

- $e = e_0.f$

- $e = v.f$

  Let $v = \langle \ell, S \rangle$. There are also two cases for the derivation $\lfloor \sigma, H, R \rfloor \vdash v.f : T, \Gamma$: T-FIN and T-GET.

  * T-FIN

    Then $T = \mathsf{ptype}(\lfloor \sigma, H, R \rfloor, v.f)$, where f is a final field. The evaluation is $v.f, \sigma, H, R \longrightarrow v', \sigma, H, R'$, where $S = P!\backslash \overline{f'}$ and $H(\ell, \mathsf{fclass}(P, f), f) = v'$. According to Lemma 5.5, $\lfloor \sigma, H, R \rfloor \vdash \mathsf{ptype}(\lfloor \sigma, H, R \rfloor, v.f) \approx \mathsf{ptype}(\lfloor \sigma, H, R \rfloor, v')$. Then by T-FIN and T-SUB, $\lfloor \sigma, H, R \rfloor \vdash v' : \mathsf{ptype}(\lfloor \sigma, H, R \rfloor, v.f), \lfloor \sigma, H, R \rfloor$.

  * T-GET

    Let $\mathsf{ftype}(\lfloor \sigma, H, R \rfloor, S, f) = T_f$. By T-GET, $T = T_f$, and $\Gamma = \lfloor \sigma, H, R \rfloor$ since $e = v.f$. According to R-GET and the definition of the auxiliary function view, $\lfloor \sigma, H, R \rfloor \vdash v' : T, \lfloor \sigma, H, R \rfloor$. The applicability of view is based on the well-formedness of the heap.

- $e = e_0.f$ where $e_0 \neq x$ and $e_0 \neq v$

  Then R-CONG is the only rule that can apply, and $e_0, \sigma, H, R \longrightarrow e_0', \sigma', H', R'$. There are also two cases for the derivation of $\lfloor \sigma, H, R \rfloor \vdash e_0.f : T, \Gamma$.

  * T-FIN

    The proof is based on Lemma 5.5, which is similar to the corresponding case in the proof for $e = v.f$.

  * T-GET

    Then $\lfloor \sigma, H, R \rfloor \vdash e_0 : T_0, \Gamma$ and $\mathsf{ftype}(\lfloor \sigma, H, R \rfloor, T_0, f) = T_f = T$. By the induction hypothesis, $\lfloor \sigma', H', R' \rfloor \vdash e_0' : T_0, \Gamma'$, where Γ' is an extension of Γ. According to Lemma 5.2, $\mathsf{ftype}(\lfloor \sigma', H', R' \rfloor, T_0, f) = T_f$. Thus by T-GET, $\lfloor \sigma', H', R' \rfloor \vdash e_0'.f : T_f, \Gamma'$.

- $e = (x.f = e_0)$

  - $e = (x.f = v)$

    Then $x.f = v, \sigma, H, R \longrightarrow v, \sigma', H', R'$ where $\sigma' = \mathsf{grant}(\sigma, x.f)$, $\sigma(x) = \langle \ell, S \rangle$, $H' = H[\langle \ell, \mathsf{fclass}(P, f), f \rangle := v]$ where $S = P!\backslash \overline{f'}$, and $R' = R, \langle P!\backslash(\overline{f'} - f) \rangle$. Then T-SET applies to the derivation of $\lfloor \sigma, H, R \rfloor \vdash e : T, \Gamma$: $\Gamma = \mathsf{grant}(\lfloor \sigma, H, R \rfloor, x.f)$, and $\lfloor \sigma', H', R \rfloor = \mathsf{grant}(\lfloor \sigma, H, R \rfloor, x.f)$, so $\lfloor \sigma', H', R \rfloor$ is a trivial extension of Γ. According to the definition of runtime typing environment, $\lfloor \sigma', H', R' \rfloor$ may add to $\lfloor \sigma', H', R \rfloor$ several path equivalence relationships, and therefore an extension of $\lfloor \sigma', H', R \rfloor$, which means that $\lfloor \sigma', H', R' \rfloor$ is an extension of Γ.

  - $e = (x.f = e_0)$ where $e_0 \neq v$

    Then R-CONG can apply, and the induction hypothesis can be directly used.

- $e = (\texttt{final}\ S\ x = e_1; e_2)$

  - $e = (\texttt{final}\ S\ x = v; e_2)$

    R-LET is the only rule that can apply: $e, \sigma, H, R \longrightarrow e_2, \sigma', H, R$ where $\sigma' = \sigma[x := v]$ and x is fresh variable. By T-LET and T-FIN, $\lfloor \sigma, H, R \rfloor \vdash v : S, \lfloor \sigma, H, R \rfloor$ and $\lfloor \sigma, H, R \rfloor, x : S \vdash e_2 : T, \Gamma_2$, and $\Gamma_2 = \Gamma, x : T'$ for some T'. Since $\sigma' = \sigma[x := v]$ and x is fresh, we have $\lfloor \sigma', H, R \rfloor = \lfloor \sigma, H, R \rfloor, x : T''$ for some T'', where $\vdash T'' \leq S$, according to the definition of the run-time typing environment. Therefore $\lfloor \sigma', H, R \rfloor \vdash e_2 : T, \Gamma_2'$, where $\Gamma_2' = \Gamma', x : T'''$ for some T'''. It is then easy to see that Γ' is an extension of Γ.

  - $e = (\texttt{final}\ S\ x = e_1; e_2)$ where $e_1 \neq v$

    Then R-CONG can apply, and therefore the induction hypothesis can be used.

- $e = e_0.m(\overline{e})$

- $e = \langle \ell, S \rangle.m(\overline{v})$

  Then R-CALL is the only rule that can apply: $\langle \ell, S \rangle.m(\overline{v}), \sigma, H, R \longrightarrow e_m\{y_0/\mathtt{this}, \overline{y}/\overline{x}\}, \sigma', H, R$, where $\sigma' = \sigma[y_0 := \langle \ell, S \rangle, \overline{y} := \overline{v}]$, and $y_0, y_1, \ldots, y_n \notin \mathrm{dom}(\sigma)$. Therefore, $\lfloor \sigma', H, R \rfloor$ is an extension of $\lfloor \sigma, H, R \rfloor$. According to M-OK and simple α renaming, $\lfloor \sigma, H, R \rfloor, y_0 : T_0, \overline{y} : \overline{T} \vdash e_m\{y_0/\mathtt{this}, \overline{y}/\overline{x}\} : T_{n+1}$, where $\mathrm{mbody}(S, m) = T_{n+1}\, m(\overline{T}\, \overline{x})\, \{e_m\}$. Now applying Lemma 5.3 and Lemma 5.4, for $y_0, y_1, \ldots, y_n$, we can get $\lfloor \sigma', H, R \rfloor \vdash e_m : T_{n+1}^{n+1}$, where $T_i^0 = T_i$ and $T_i^j = T_i^{j-1}\{\lfloor \sigma', H, R \rfloor; T_{j-1}^{j-1}/y_{j-1}\}$. $T_{n+1}^{n+1} = T$ in T-CALL. Let $\lfloor \sigma, H, R \rfloor \vdash e : T, \Gamma$ and $\lfloor \sigma', H, R \rfloor \vdash e_m\{y_0/\mathtt{this}, \overline{y}/\overline{x}\} : T, \Gamma'$. It is obvious that Γ' is an extension of Γ, because $\mathrm{FV}(e_m) \subseteq \{\mathtt{this}, \overline{x}\}$.

- $e = e_0.m(\overline{e})$ where $e_i \neq v$ for some $i \in \{0, 1, \ldots, n\}$

  Then R-CONG can apply. This case is easy to prove by the induction hypothesis.

- $e = \mathtt{new}\ S$

  Then R-ALLOC can apply: $\mathtt{new}\ S, \sigma, H, R \longrightarrow e', \sigma, H, R'$, where $e' = \mathtt{final}\ S! \backslash \overline{f}\ x = v;\ x.\overline{f} = \overline{e}\{x/\mathtt{this}\}$; x, and $\mathrm{fields}(S) = [\mathtt{final}]\ \overline{T_f}\ \overline{f} = \overline{e}$, and $v = \langle \ell, S! \backslash \overline{f} \rangle$, and $R' = R, v$. It is easy to see that $\lfloor \sigma, H, R' \rfloor \vdash e' : S!$, because all the masks are going to be removed by the field initializers. It is also obvious that $\lfloor \sigma, H, R' \rfloor$ is an extension of $\lfloor \sigma, H, R \rfloor$ by the definition of run-time typing environments.

- $e = e_1;\ e_2$

  - $e = v;\ e_2$

    Then R-SEQ can apply: $e, \sigma, H, R \longrightarrow e_2, \sigma, H, R$. By T-SEQ, $\lfloor \sigma, H, R \rfloor \vdash v : T_1, \Gamma_1$ and $\Gamma_1 \vdash e_2 : T, \Gamma$. Also, it is easy to see that $\Gamma_1 = \lfloor \sigma, H, R \rfloor$. Therefore $\lfloor \sigma, H, R \rfloor \vdash e_2 : T, \Gamma$.

  - $e = e_1;\ e_2$ where $e_1 \neq v$

    Then R-CONG can apply, and the proof is by the induction hypothesis.

- $e = (\mathtt{view}\ S)e_0$

  - $e = (\mathtt{view}\ S)\langle \ell, S' \rangle$

    Then R-VIEW is the only rule that can apply. By T-VIEW, $\lfloor \sigma, H, R \rfloor \vdash (\mathtt{view}\ S)\langle \ell, S' \rangle : S$, and there is a type T such that $\lfloor \sigma, H, R \rfloor \vdash \langle \ell, S' \rangle : T$ and $\lfloor \sigma, H, R \rfloor \vdash T \rightsquigarrow S$. According to the definition of view, $\lfloor \sigma, H, R \rfloor \vdash \mathrm{view}(\langle \ell, S' \rangle, S) : S$, and since $\lfloor \sigma, H, R' \rfloor$, where $R' = R, \mathrm{view}(\langle \ell, S' \rangle, S)$, is an extension of $\lfloor \sigma, H, R \rfloor$, we have $\lfloor \sigma, H, R' \rfloor \vdash \mathrm{view}(\langle \ell, S' \rangle, S) : S$.

    Now it remains to prove that the application of the auxiliary function view is well-defined. If $S' \leq S$, it is obviously well-defined. Otherwise, let $S' = P'! \backslash \overline{f'}$, and $S = PS \backslash \overline{f}$, and we need to prove that $\exists! P\ .\ \exists \overline{f''} \supseteq \overline{f'}\ .\ \vdash P! \leq PS \wedge \vdash P'! \backslash \overline{f''} \rightsquigarrow P! \backslash \overline{f}$.

    As shown above, there is a type T such that $\lfloor \sigma, H, R \rfloor \vdash T \rightsquigarrow S$, and therefore $\emptyset \vdash T \rightsquigarrow S$, according to the definition of $\lfloor \sigma, H, R \rfloor$. Then T has to be non-dependent. Also, it is obvious that $\vdash P'! \backslash \overline{f'} \leq T$, so $\overline{f'}$ must be a subset of the masks in T, which is $\overline{f''}$. Now the proof goes on by induction on the derivation of $\emptyset \vdash T \rightsquigarrow S$:

    * SH-REFL

      Vacuously true, since it implies that $\vdash S' \leq S$, which is the first case of the definition of view.

    * SH-TRANS

      By the inner induction hypothesis.

    * SH-ENV

      Vacuously true.

    * SH-MASK

      By the inner induction hypothesis.

* SH-DECL

    Then both T and S are exact types, and therefore $S = P!\backslash\overline{f}$ and $T = P'!\backslash\overline{f''}$, by the definition of exact types.

* SH-CLS

    The premise of SH-CLS is exactly what we want to prove.

– $e = (\text{view } S)e_0$ where $e_0 \neq v$

    Then R-CONG can apply. This case can be easily proved using the induction hypothesis.

$\square$

Lemma 5.7 *(Progress) If $\vdash e, \sigma, H, R$ and $\lfloor \sigma, H, R \rfloor \vdash e{:}T$, then $e = v$, or there is a configuration e', σ', H', R' such that $e, \sigma, H, R \longrightarrow e', \sigma', H', R'$.*

PROOF: The proof is by structural induction on e.

* $e = v$

    Trivial.

* $e = x$

    Then $\lfloor \sigma, H, R \rfloor \vdash x{:}T$, and by the definition of $\lfloor \sigma, H, R \rfloor$, $\sigma(x) = \langle \ell, T \rangle$. Therefore R-VAR applies, and we have $e' = \langle \ell, T \rangle$, $\sigma' = \sigma$, $H' = H$, and $R' = R$.

* $e = e_0.f$

    If $e_0 \neq v$, then it is easy to see that R-CONG applies, and the evaluation can make progress.

    If $e_0 = v$ where $v = \langle \ell, P!\backslash\overline{f'} \rangle$, then by T-GET, and the definition of ftype, we know $f \notin \overline{f'}$. Then according to configuration well-formedness CONFIG, $H(\ell, \text{fclass}(P), f) = v'$, and $\text{view}(v', S)$ is well-defined ($S = \text{ftype}(\emptyset, P!\backslash\overline{f'}, f)$). Therefore R-GET applies, and the evaluation can make progress.

* $e = x.f = e_0$

    If $e_0 \neq v$, then R-CONG applies.

    Suppose $e_0 = v$. By T-SET and T-FIN, $\lfloor \sigma, H, R \rfloor \vdash v{:}T_v$, $\lfloor \sigma, H, R \rfloor$, and then by T-SET, $\lfloor \sigma, H, R \rfloor \vdash x{:}T_x$. By the definition of $\lfloor \sigma, H, R \rfloor$, we have $\sigma(x) = \langle \ell, T_x \rangle$. It must be the case that $T_x = P!\backslash\overline{f'}$, for some P and $\overline{f'}$, because the view T_x is always a non-dependent exact type. Then R-SET applies, and $e' = v$, $\sigma' = \text{grant}(\sigma, x.f)$, $H' = H[\langle \ell, \text{fclass}(P, f), f \rangle := v]$, and $R' = R$.

* $e = e_0.m(\overline{e})$

    If e_0 or any e_i of $\overline{e}$ is not a value, R-CONG applies, and the evaluation can make progress.

    Otherwise, $e = v_0.m(\overline{v})$, where $v_0 = \langle \ell_0, S_0 \rangle$. By T-CALL, and the definitions of mtype and mbody, method lookup succeeds and we have $\text{mbody}(S, m) = T_{n+1} \ m(\overline{T} \ \overline{x}) \ \{e_m\}$. Then R-CALL applies, and $e' = e_m\{y_0/\texttt{this}, \overline{y}/\overline{x}\}$ where all the y_i are fresh variables, $\sigma' = \sigma[y_0 := \langle \ell_0, S_0 \rangle, \overline{y} = \overline{v}]$, $H' = H$, and $R' = R$.

* $e = e_1; e_2$

    If $e_1 = v$, then R-SEQ applies, and $e' = e_2$, $\sigma' = \sigma$, $H' = H$, and $R' = R$; otherwise, by the induction hypothesis, there exists e_1', σ', H', and R', such that $e_1, \sigma, H, R \longrightarrow e_1', \sigma', H', R'$. Then R-CONG applies.

- $e = \text{new } T$

  If T is not a non-dependent type, then the type T can be further evaluated.

  Now suppose $T = S$. By R-ALLOC, $e' = \text{final } S!\backslash \overline{f}\ x = v;\ x.\overline{f} = \overline{e}\{x/\text{this}\};\ x$, where $\overline{f}$ is the set of all fields of S, x is a fresh variable, and $v = \langle \ell, S!\backslash \overline{f} \rangle$ with a fresh location ℓ. We also have $\sigma' = \sigma$, $H' = H$, and $R' = R, v$.

- $e = (\text{view } T)e_0$

  If $e_0 \neq v$ or $T \neq S$, then R-CONG can apply.

  If $e_0 = v$ and $T = S$ where S is a non-dependent type, then $e' = \text{view}(v, S)$, $\sigma' = \sigma$, $H' = H$, and $R' = R, \text{view}(v, S)$. We only need to prove that $\text{view}(v, S)$ is well-defined. Let $S = PS\backslash \overline{f}$ and $v = \langle \ell, P'!\backslash \overline{f'} \rangle$. By T-FIN and T-VIEW, $\lfloor \sigma, H, R \rfloor \vdash P'!\backslash \overline{f''} \rightsquigarrow PS\backslash \overline{f}$ and $\overline{f'} \subseteq \overline{f''}$. Now the proof is by induction on the derivation of $\lfloor \sigma, H, R \rfloor \vdash P'!\backslash \overline{f''} \rightsquigarrow PS\backslash \overline{f}$:

  - SH-REFL

    Then $P'!\backslash \overline{f''} = PS\backslash \overline{f}$, and $\text{view}(v, S)$ is trivially well-defined.

  - SH-TRANS

    Then there exists a non-dependent type T', such that $\lfloor \sigma, H, R \rfloor \vdash P'!\backslash \overline{f''} \rightsquigarrow T'$ and $\lfloor \sigma, H, R \rfloor \vdash T' \rightsquigarrow S$. Note that $\lfloor \sigma, H, R \rfloor$ contains no sharing relationships, which means that the T' has to be non-dependent, and $\lfloor \sigma, H, R \rfloor$ can actually be replaced with empty environment in the above two judgments. By the induction hypothesis, $\text{view}(v, T') = \langle \ell, P''!\backslash \overline{f'''} \rangle$ is well-formed, and $\text{view}(\langle \ell, P''!\backslash \overline{f'''} \rangle, S)$ is also well-formed. Now just observe that any combination of the two cases in the definition of the auxiliary function view implies the well-formedness of $\text{view}(v, S)$.

  - SH-ENV

    Vacuously true, because $\lfloor \sigma, H, R \rfloor$ does not contain any sharing relationships.

  - SH-MASK

    It is easy to see that adding a mask on both types does not affect the applicability of view.

  - SH-DECL

    Then S is an exact type, and is itself the unique target type in the result of the view function.

  - SH-CLS

    The premise directly implies that $\text{view}(v, S)$ is well-defined.

- $e = \text{final } T\ x = e_1;\ e_2$

  If $e_1 \neq v$ for any value v, then by T-LET, $\lfloor \sigma, H, R \rfloor \vdash e_1 : T_1, \Gamma$. Therefore, by the induction hypothesis, there exists e_1', σ', H', and R', such that $e_1, \sigma, H, R \longrightarrow e_1', \sigma', H', R'$. Then R-CONG applies.

  If $e_1 = v$, then R-LET applies, and $e' = e_2\{y/x\}$ where y is a fresh variable, $\sigma' = \sigma[y := v]$, $H' = H$, and $R' = R$.

$\square$

Now we can prove the soundness theorem of the J&$_s$ calculus.

Theorem 5.8 (*Soundness*) *If $\vdash \langle \overline{L}, e \rangle$ ok, and $\emptyset \vdash e : T$, and $e, \emptyset, \emptyset, \emptyset \rightarrow^* e', \sigma, H, R$, then either $e' = v$ and $\lfloor \sigma, H, R \rfloor \vdash v : T$, or $\exists e'', \sigma', H', R'\ .\ e', \sigma, H, R \longmapsto e'', \sigma', H', R'$.*

PROOF: Follows from Lemma 5.6 and Lemma 5.7. $\square$

6 Implementation

We have implemented a prototype compiler of J&$_s$ in the Polyglot framework [30]. The compiler is an extension based on the latest implementation of the Jx/J& compiler [29, 31], which is itself an extension of the Polyglot base Java compiler. The J&$_s$ compiler extension has 5,200 lines of code, excluding blank lines and comments. J&$_s$ also has a run-time system that supports implicit classes and provides machinery for view changes and run-time type inspection. The run-time system is written in Java, and has 1,500 lines of code, most of which implements a custom classloader, using the ASM bytecode manipulation framework [8].

J&$_s$ is implemented as a translation to Java. The amount of code produced by the compiler is proportional to the size of the source code. The compiler generates class declarations only for explicit classes. For implicit classes, the custom classloader lazily synthesizes classes representing them at run time.

6.1 Type checking

Because of masks, J&$_s$ has a flow-sensitive type system. Every method is checked in two phases. The first phase is flow-insensitive, ignoring masks, and generates typing information necessary for building the control-flow graph. The second phase is essentially an intraprocedural data-flow analysis, which computes a type binding, possibly with masks, for each local variable (including `this`) at every program point.

6.2 Translating classes

The translation is similar to but more efficient than the J& translation described in [31], primarily because of the custom classloader.

First, masks are erased after type checking. Then, every explicit J&$_s$ class is translated into several Java classes and interfaces, among which the most important ones are the *instance class* and the *class class*. The instance class contains all the object fields of the J&$_s$ class. At run time, each object of a J&$_s$ class is represented as an object of the instance class. The class class contains method implementations, static fields, and type information needed to implement casts, `instanceof`, and prefix types. The class class also includes getter and setter methods for all object fields, since unshared fields may be duplicated, and therefore, field accesses are also view-dependent. Various interfaces are generated for simulating multiple inheritance.

An implicit class also has an instance class and a class class, both synthesized by the custom classloader. The instance class collects all object fields from superclasses. The class class contains a dispatch method for each inherited J&$_s$ method. The dispatch method calls the appropriate method implementation in the class class of some explicit J&$_s$ class, identified by the classloader.

For each set of shared J&$_s$ classes, the classloader maintains a representative instance class that collects all the object fields, including the duplicate ones. At run time, objects of all these shared classes are created as instances of the representative instance class. When a new shared class that contains more object fields is loaded, the classloader updates the representative instance class to include the new fields. All the existing objects of the old representative instance class will be lazily converted to the most up-to-date representative instance class. The conversion works because objects of instance classes are referenced indirectly, as described in Section 6.3.

6.3 Supporting views and view changes

A J&$_s$ object can be an instance of several shared classes, and every reference to the object could have a different view. Therefore, the J&$_s$ implementation adds a level of indirection. Each object is referenced indirectly through a *reference object*, containing two fields: one points to the instance class object; and the other points to a class class object, which is the view associated with the reference.

29

The view determines the behavior of the translated J&$_s$ object. Method calls are dispatched on the view rather than on the object itself. Prefix types are evaluated with the view. Casts and `instanceof`s check the view. Which copy of a duplicated field is accessed also depends on the view.

A view change operation $(\texttt{view}\ T)e$ is translated to generating a new reference object with the same instance class object, and a view compatible with T. The implementation memoizes the result to the most recent view change operation on any reference object, to avoid repeatedly generating the same reference object.

A J&$_s$ object might obtain a new view implicitly. Moving an object from one family to another would implicitly move all the other objects that are transitively reachable through shared fields. Implicit view changes are carried out lazily, only at the time when objects are accessed through fields.

Reference objects add some overhead for accessing members of an object. With a lower-level target language, a probably faster implementation is possible in which different object views are represented as different pointers into the same object, like the C++ implementation of multiple inheritance. Since methods are dispatched on views, the object would contain pointers to several dispatch tables, one for each view. Method dispatch should then have performance similar to C++ virtual method calls.

6.4   Java compatibility

J&$_s$ is mostly backward-compatible with Java. Any Java code with no nested classes is also legal J&$_s$ code. J&$_s$ programs may use existing Java code, including libraries that are compiled by a Java compiler. Of course, precompiled Java code does not enjoy the benefit of class sharing. The J&$_s$ reference object forwards method calls to `hashCode`, `equals`, etc., and therefore J&$_s$ objects may be passed to Java code, for example, to be stored in a `HashSet`. There is one limitation in the current implementation: a J&$_s$ class cannot be shared if it inherits from any Java superclass other than `java.lang.Object`, or if it implements most Java interfaces. This is the artifact of using reference objects to support views; a lower-level implementation should not have this limitation.

The J&$_s$ language currently extends Java 1.4. It does not support generics, which seem to be an orthogonal feature. We leave the interaction between generics and class sharing to future work.

6.5   Concurrency

Class sharing is compatible with multi-threaded code. The J&$_s$ compiler ensures that in the generated code, `synchronized` statements and methods are synchronized on the wrapped instance class object, rather than on the reference object. The implementation of the reference object also relays method calls to `wait`, `notify` and `notifyAll` to the underlying instance class object. A view change operation usually just creates a new reference object that contains a reference to an existing instance class object, without affecting existing references, and therefore requires no synchronization. The only operation that needs to be synchronized is the lazy conversion of an object to the most up-to-date instance class.

7   Experience

7.1   Jolden benchmarks

We tested the J&$_s$ implementation with the jolden benchmarks [9] to study the performance overhead for code that does not use the new extensibility features of J&$_s$. All ten benchmarks, with few changes, are tested with four language implementations: Java, J& as described in [31], J& with a classloader similar to that described in Section 6.2, and J&$_s$. Table 1 compares the results. The testing hardware is a Lenovo

	bh	bisort	em3d	health	mst	perimeter	power	treeadd	tsp	voronoi
Java	1.74	0.53	0.17	0.41	0.88	0.22	1.00	0.14	0.12	0.31
J& [31]	13.91	1.77	0.48	8.45	4.43	2.82	2.43	2.20	0.37	7.19
J& with classloader	2.02	0.71	0.22	0.71	1.06	0.39	1.14	0.21	0.16	0.59
J&$_s$	2.61	0.88	0.23	1.61	1.54	0.47	1.27	0.45	0.17	0.83

Table 1: Results for the jolden benchmarks. Average time over ten runs, in seconds.

Height	16	18	20
Tree creation	0.110	0.287	1.295
Traversal before view changes	0.008	0.027	0.105
View changes	0.125	0.367	1.303
Traversal after view changes	0.006	0.025	0.099
Explicit translation	0.145	0.622	1.669

Table 2: Results for tree traversal. Average time over ten runs, in seconds.

Thinkpad T60 with Intel T2600 CPU and 2GB memory, and the software environment consists of Microsoft Windows XP, Cygwin, and JVM 1.6.0_07.

Table 1 shows that the use of a custom classloader greatly improves the performance, comparing the two implementations of J&. Unsurprisingly, nested inheritance and class sharing do introduce overhead. The J&$_s$ times show a 37% slowdown versus the classloader-based J& implementation, and 94% versus the highly optimized Java HotSpot VM. The overhead seems reasonable, especially considering that the current implementation works as a source-to-source translation to Java, precluding many optimizations and implementation techniques. We expect that a more sophisticated implementation could remove much of the overhead.

Running programs in J&$_s$ have the latent capability to be extended in many ways, so it is not surprising that there is some performance cost. But it seems software designers are often willing to pay a cost for extensibility, because they often add indirections and use design patterns that promote extensibility but have run-time overhead. We believe that J&$_s$ removes the need for many such explicit extensibility hooks, while making code simpler. For systems where extensibility is more important than high performance [46], the existing implementation may already be fast enough.

7.2 Tree traversal

The jolden benchmarks do not use the new features provided by J&$_s$. To study the performance of view changes, especially on large data structures, we wrote a small benchmark in which two families share classes that implement binary trees. A complete binary tree of a given height is first created in the base family, and an explicit view change is applied to the root of the tree. A depth-first traversal is carried out to trigger all the lazy implicit view changes. The testing environment is the same as in Section 7.1.

Table 2 summarizes the results. In-place adaptation, even with a traversal that triggers all the implicit view changes, is faster than an explicit translation that creates new objects in the derived family, and the running time is also close to that of the initial creation of the tree. The fourth row shows that once the implicit view changes are complete and the new reference objects have been memoized, traversals execute as fast as a traversal before the view changes. This benchmark does a complete traversal. Since view changes are lazily triggered, the relative performance of adaptation would look even better if not all nodes needed to be visited after adaptation.

Figure 20: Lambda compiler structure. Translator and some AST nodes not shown.

7.3 Lambda compiler

The J&$_s$ implementation successfully compiles and executes the completed version of the lambda compiler shown in Figure 6, which compiles λ-calculus enhanced with sums and pairs down to simple λ-calculus.

The structure of the lambda compiler is illustrated in Figure 20. Solid arrows represent class and family inheritance, and dashed arrows represent sharing declarations.

The lambda compiler has a `base` family with classes representing AST nodes for simple λ-calculus. There are two families directly derived from the base family, extended with sums and pairs respectively. The `sum` family and the `pair` family each share AST classes from the `base` family, and each implement in-place translation to simple λ-calculus, as shown in Figure 6. The last derived family `sumpair` composes the `sum` family and the `pair` family, leading to a compiler that supports both sums and pairs. The `sumpair` family shares `Abs` and other AST classes from the simple λ-calculus with the `base` family, and by transitivity, with `sum` and `pair`. The code of `sumpair` just sets up the sharing relationships, without a single line of translation code. The in-place translation code from `sum` and `pair` is composed to translate away sums and pairs.

This program is about 250 lines long, but uses the features of J&$_s$ in a sophisticated way. For example, families have both shared and unshared classes, masked types are used to ensure objects of new AST classes (pair and sum) are translated away, and the two translations are composed to translate pairs and sums at once. The lambda compiler example is inspired by the Polyglot framework, and it encapsulates most of the interesting issues that arise in making Polyglot extensible. The results suggest Polyglot would be simpler in J&$_s$, but this is left for future work.

7.4 CorONA

Our second significant example shows that the adaptation capability of J&$_s$ can be used to seamlessly upgrade a running server and its existing state with entirely new functionality. We ported CorONA, an RSS feed aggregation system [37], to J&$_s$, and successfully used class sharing to update the system at run time with a new caching algorithm.

CorONA is originally an extension of Beehive [38], which is itself an extension of Pastry [40] that provides a distributed hash table. Beehive extends Pastry with active replication, whereas PC-Pastry [38] extends Pastry with passive caching. We refactored the ported CorONA, and composed it with PC-Pastry and Beehive respectively, creating two applications (named PCCorONA and BeeCorONA) with different caching strategies. The system is tested by first running PCCorONA for a while, and then evolving the running system from passive caching to active caching by compiling and loading a new package, BeeCorONA.

Sharing declarations are added so that classes representing host nodes, data objects, and network addresses are shared between CorONA and its two derived families. Classes for network messages and cache

management are not shared. The evolution code goes over all the host nodes, which are the top-level objects in the system, changing their views, and creating new caching managers.

The amount of code to implement evolution is relatively small (less than 40 lines of code, compared to 8300 for the whole system, excluding comments and empty lines). Very little code is needed because we only need to change the views of host nodes; all the other referenced objects will have their views changed implicitly when they are accessed. In the host node class, fields storing the caching managers are not shared, and masked types ensure that they are initialized in the evolved system. With a slightly different configuration, we can actually run the two variants of the system at the same time, using the same set of host node objects. View-dependent types ensure that network messages of correct versions are created and accepted for each of the systems.

8   Related work

Adaptation.   The Adapter design pattern [19] is a protocol for implementing adaptation. However, this and other related patterns are tedious and error-prone to implement, rely on statically unsafe type casts, and do not preserve object identity or provide bidirectional adaptation as J&$_s$ does.

Expanders [50] are a mechanism for adaptation. New fields, methods, and superinterfaces can be added into existing classes. Expanders are more expressive than *open classes* [11], which can only add new methods. Method dispatch is statically scoped, so expanders do not change the behavior of existing clients. New state is added by wrapper classes; a map from objects to wrappers ensures uniqueness of wrapper instances.

CaesarJ [27, 1] is an aspect-oriented language that supports adaptation with *wrappers* called *aspect binders*. Wrappers and expanders are similar. They both can extend wrapped classes with new states, operations, and superinterfaces; no duplicate wrappers are created for objects; and dynamic wrapper selection is similar to expander overriding. Wrappers in CaesarJ are less transparent: a wrapper constructor must be called to get a wrapper instance, whereas expander operations can be applied directly to objects. Multiple inheritance in CaesarJ makes wrapper selection ambiguous; J&$_s$ disambiguates via views.

Both expanders and CaesarJ wrappers share limitations: it is impossible to override methods in the original family, and therefore there is no dynamic dispatching across families; object identity is not preserved; the use of expanders and wrappers is limited to adaptation, since the adapter family cannot be used independently of its original family. Class sharing in J&$_s$ provides more flexibility.

The Fickle$_{III}$ [14] language has a *re-classification* operation for objects to change their classes. Re-classifications are similar to view changes in J&$_s$, but directly change the behavior of all existing references to the object; therefore, effects are needed to track the change. A re-classification might leave fields uninitialized, while masked types in J&$_s$ ensure that after a view change, fields are initialized before use. Fickle$_{III}$ does not support class families.

Chai$_3$ [42] allows traits to be dynamically substituted to change object behaviors, similar to view changes in J&$_s$. However, *Chai$_3$* does not support families, and the fact that traits do not have fields makes it harder to support manipulation of data structures.

Some work on adaptation, including pluggable composite adapters [28], object teams [20], and delegation layers [35], also has some notion of families of classes. However, family extensibility does not apply to the relationship between the family of adapter classes and the family of adaptee classes. Therefore, these mechanisms either do not support method overriding and dynamic dispatch between the adapter and adaptee families [28, 20], or have a weaker notion of families in which programmers have to manually "wire" inheritance relationships between the base family and the delegation family [35]. These mechanisms all use lifting and lowering, introduced in [28], to convert between adapter and adaptee classes. Lifting and lowering are similar to view changes in J&$_s$, but are not symmetric and do not support late binding.

33

Family inheritance.   Several different mechanisms have been proposed to support family inheritance, including virtual classes, nested inheritance, variant path types, mixin layers, etc. In all these family inheritance mechanisms, families of classes are disjoint, whereas with class sharing, different families can share classes and their instances.

Virtual classes [26, 25, 16, 18, 10] are inner classes that can be overridden just like methods. Path-dependent types are used to ensure type safety. The soundness of virtual classes has been formally proved by Ernst et al. [18], and by Clarke et al. [10].

Nested inheritance [29] supports overriding of nested classes, which are similar to virtual classes. Nested intersection [31] adds and generalizes intersection types [39, 12] in the context of nested inheritance to provide the ability to compose extensions.

Virtual classes support *family polymorphism* [17], where families are identified by the enclosing instance. Nested inheritance supports what Clarke et al. [10] called *class-based family polymorphism*, where each dependent class defines a family of classes nested within and also enclosing it. With prefix types, any instance of a class in the family can be used to name the family.

Variant path types [23] support family inheritance without dependent types, using exact types and relative path types (similar to `this.class`) for type safety. These exact types are very different from those in J&$_s$: in a J&$_s$ exact type $A.B!.C$, exactness applies to the whole prefix before !, that is, $A.B$; in an exact type $A@B.C$ in [23], exactness applies to the simple type name right after @, that is, B.

Mixin layers [41] generalizes *mixins* [4]. Mixins are classes that can be instantiated with different super-classes, and mixin layers are mixins that encapsulate other mixins. Mixin layers support family inheritance: when a mixin layer is instantiated, all the inner mixins are instantiated correspondingly.

Virtual types [47, 7, 48, 21] are type declarations that can be overridden. Virtual types are more limited than virtual classes: they provide family polymorphism but not family inheritance. Scala [33, 34] supports family polymorphism and composition through virtual types, path-dependent types, and mixin composition. It also supports parametric polymorphism. Scala does not have virtual classes and does not support family inheritance. Scala has *views* that are implicitly-called conversion functions. Scala views do not provide adaptation: they create new instances in the target types.

Sharing in functional languages.   Wadler's views [49] are an isomorphism between a new data type and an existing one, which is similar to a sharing declaration. Of course, there are obvious differences: Wadler's views are for a functional setting, and primarily relate abstract data types and types inductively defined with pattern matching, whereas J&$_s$ creates sharing relationships only between overridden and overriding classes.

The SML module system [24] has *sharing constraints*, which require functor module parameters to agree on type components. SML sharing constraints dictate applicability of functors; J&$_s$ sharing constraints dictate applicability of method bodies.

Safe dynamic software updating.   There is prior work on the problem of safely updating software without downtime.

Barr and Eisenbach [2] propose a framework to support dynamic update of Java components that satisfy the binary compatibility requirement [45], which also includes a custom classloader. The goal is to provide a tool that keeps Java libraries up to date, rather than to improve the extensibility of the language.

Duggan [15] describes a calculus that combines Wadler's views and SML sharing constraints to support hot-swapping modules. J&$_s$ differs because it does not copy values across different versions; instead, it generates different views on the same object.

Proteus [43] finds proper timing for a given global update with static analysis. Abstract and concrete types in Proteus bear some resemblance to inexact and exact types in J&$_s$: abstractly typed variables allow values of different concrete types, and inexactly typed variables may store objects with different exact

views.

9 Conclusions

This paper introduces class sharing, a flexible extensibility mechanism that combines the advantages of both family inheritance and adaptation. As several examples show, class sharing adds new capabilities such as family adaptation, dynamic object evolution, and in-place translation. These capabilities are supported by a variety of mechanisms. Dynamic views typed with dependent classes statically track the families of values; masked types ensure shared and unshared classes can be mixed safely. The language is proved sound, showing that the new extensibility provided by class sharing does not come at a price in type safety.

Acknowledgments

References

[1] Ivica Aracic, Vaidas Gasiunas, Mira Mezini, and Klaus Ostermann. An overview of CaesarJ. In Awais Rashid and Mehmet Aksit, editors, *Lecture Notes in Computer Science : Transactions on Aspect-Oriented Software Development I*, pages 135–173. Springer-Verlag, 2006.

[2] Miles Barr and Susan Eisenbach. Safe upgrading without restarting. In *Proceedings of 19th International Conference on Software Maintenance (ICSM)*, pages 129–137, 2003.

[3] Alexandre Bergel, Stéphane Ducasse, and Oscar Nierstrasz. Classbox/J: Controlling the scope of change in Java. In *Proc. 20th ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 177–189, San Diego, CA, USA, October 2005.

[4] Gilad Bracha and William Cook. Mixin-based inheritance. In Norman Meyrowitz, editor, *Proc. 5th ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 303–311, Ottawa, Canada, 1990. ACM Press.

[5] Frederick P. Brooks. *The Mythical Man-Month: Essays on Software Engineering*. Addison-Wesley, 1978.

[6] Kim B. Bruce. Safe static type checking with systems of mutually recursive classes and inheritance. Technical report, Williams College, 1997. `http://cs.williams.edu/~kim/ftp/RecJava.ps.gz`.

[7] Kim B. Bruce, Martin Odersky, and Philip Wadler. A statically safe alternative to virtual types. In *European Conference on Object-Oriented Programming (ECOOP)*, pages 523–549, Brussels, Belgium, July 1998.

[8] Eric Bruneton, Romain Lenglet, and Thierry Coupaye. ASM: A code manipulation tool to implement adaptable systems, 2002. `http://asm.objectweb.org/current/asm-eng.pdf`.

[9] B. Cahoon and K. S. McKinley. Data flow analysis for software prefetching linked data structures in Java. In *International Conference on Parallel Architectures and Compilation Techniques (PACT)*, September 2001.

[10] Dave Clarke, Sophia Drossopoulou, James Noble, and Tobias Wrigstad. Tribe: A simple virtual class calculus. In *AOSD '07: Proceedings of the 6th International Conference on Aspect-Oriented Software Development*, pages 121–134, 2007.

[11] Curtis Clifton, Gary T. Leavens, Craig Chambers, and Todd Millstein. MultiJava: Modular open classes and symmetric multiple dispatch for Java. In *Proc. 15th ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 130–145, 2000.

[12] Adriana B. Compagnoni and Benjamin C. Pierce. Higher order intersection types and multiple inheritance. *Mathematical Structures in Computer Science*, 6(5):469–501, 1996.

[13] Bruno C. d. S. Oliveira. Modular visitor components: A practical solution to the expression families problem. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, July 2009. to appear.

[14] Ferruccio Damiani, Sophia Drossopoulou, and Paola Giannini. Refined effects for unanticipated object reclassification: Fickle$_{III}$. In *ICTCS*, pages 97–110, 2003.

[15] Dominic Duggan. Type-based hot swapping of running modules. *Acta Inf.*, 41(4):181–220, 2005.

[16] Erik Ernst. *gbeta—A Language with Virtual Attributes, Block Structure, and Propagating, Dynamic Inheritance*. PhD thesis, Department of Computer Science, University of Aarhus, Århus, Denmark, 1999.

[17] Erik Ernst. Family polymorphism. In *Proceedings of the 15th European Conference on Object-Oriented Programming (ECOOP)*, LNCS 2072, pages 303–326, 2001.

[18] Erik Ernst, Klaus Ostermann, and William R. Cook. A virtual class calculus. In *Proc. ACM Symp. on Principles of Programming Languages (POPL) '06*, pages 270–282, January 2006.

[19] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison Wesley, Reading, MA, 1994.

[20] Stephan Herrmann. Object teams: Improving modularity for crosscutting collaborations. In *Proc. Net Object Days*, 2002.

[21] Atsushi Igarashi and Benjamin Pierce. Foundations for virtual types. In *Proceedings of the 13th European Conference on Object-Oriented Programming (ECOOP)*, pages 161–185, June 1999.

[22] Atsushi Igarashi, Benjamin Pierce, and Philip Wadler. Featherweight Java: A minimal core calculus for Java and GJ. *ACM Transactions on Programming Languages and Systems*, 23(3):396–450, 2001.

[23] Atsushi Igarashi and Mirko Viroli. Variant path types for scalable extensibility. In *Proc. 22nd ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 113–132, New York, NY, USA, 2007. ACM.

[24] David MacQueen. Modules for Standard ML. In *Proc. 1984 ACM Symposium on Lisp and Functional Programming*, pages 198–204, August 1984.

[25] O. Lehrmann Madsen and B. Møller-Pedersen. Virtual classes: A powerful mechanism for object-oriented programming. In *Proc. 4th ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 397–406, October 1989.

[26] O. Lehrmann Madsen, B. Møller-Pedersen, and K. Nygaard. *Object Oriented Programming in the BETA Programming Language*. Addison-Wesley, June 1993.

[27] M. Mezini and K. Ostermann. Conquering aspects with Caesar. In *Proceedings of the 2nd International Conference on Aspect-Oriented Software Development (AOSD)*, pages 90–100, Boston, Massachusetts, March 2003.

[28] Mira Mezini, Linda Seiter, and Karl Lieberherr. Component integration with pluggable composite adapters. *Software Architectures and Component Technology*, 2000.

[29] Nathaniel Nystrom, Stephen Chong, and Andrew C. Myers. Scalable extensibility via nested inheritance. In *Proceedings of the 19th ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 99–115, October 2004.

[30] Nathaniel Nystrom, Michael R. Clarkson, and Andrew C. Myers. Polyglot: An extensible compiler framework for Java. In Görel Hedin, editor, *Compiler Construction, 12th International Conference, CC 2003*, pages 138–152, Warsaw, Poland, April 2003.

[31] Nathaniel Nystrom, Xin Qi, and Andrew C. Myers. J&: Nested intersection for scalable software composition. In *Proc. 21st ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 21–36, October 2006.

[32] Nathaniel Nystrom, Xin Qi, and Andrew C. Myers. Nested intersection for scalable software composition. Technical report, Computer Science Dept., Cornell University, September 2006. `http://www.cs.cornell.edu/nystrom/papers/jet-tr.pdf`.

[33] Martin Odersky, Philippe Altherr, Vincent Cremet, Burak Emir, Sebastian Maneth, Stéphane Micheloud, Niko-lay Mihaylov, Michel Schinz, Erik Stenman, and Matthias Zenger. An overview of the Scala programming language, June 2004. http://scala.epfl.ch/docu/files/ScalaOverview.pdf.

[34] Martin Odersky and Matthias Zenger. Scalable component abstractions. In *Proc. 20th ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 41–57, San Diego, CA, USA, October 2005.

[35] Klaus Ostermann. Dynamically composable collaborations with delegation layers. In *Proceedings of the 16th European Conference on Object-Oriented Programming (ECOOP)*, volume 2374 of *Lecture Notes in Computer Science*, pages 89–110, Málaga, Spain, 2002. Springer-Verlag.

[36] Xin Qi and Andrew C. Myers. Masked types for sound object initialization. In *Proc. ACM Symp. on Principles of Programming Languages (POPL) '09*, pages 53–65, January 2009.

[37] Venugopalan Ramasubramanian, Ryan Peterson, and Emin Gün Sirer. Corona: A high performance publish-subscribe system for the World Wide Web. In *Proceedings of Networked System Design and Implementation (NSDI)*, May 2006.

[38] Venugopalan Ramasubramanian and Emin Gün Sirer. Beehive: $O(1)$ lookup performance for power-law query distributions in peer-to-peer overlays. In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, March 2004.

[39] John C. Reynolds. Design of the programming language Forsythe. Technical Report CMU-CS-96-146, Carnegie Mellon University, June 1996.

[40] Antony Rowstron and Peter Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. In *IFIP/ACM International Conference on Distributed Systems Platforms (Middleware)*, pages 329–350, November 2001.

[41] Yannis Smaragdakis and Don Batory. Mixin layers: An object-oriented implementation technique for re-finements and collaboration-based designs. *ACM Transactions on Software Engineering and Methodology*, 11(2):215–255, April 2002.

[42] Charles Smith and Sophia Drossopoulou. Chai: Traits for Java-like languages. In *Proceedings of the 19th European Conference on Object-Oriented Programming (ECOOP)*, pages 453–478, 2005.

[43] Gareth Stoyle, Michael Hicks, Gavin Bierman, Peter Sewell, and Iulian Neamtiu. Mutatis mutandis: safe and predictable dynamic software updating. In *Proc. ACM Symp. on Principles of Programming Languages (POPL) '05*, pages 183–194, 2005.

[44] Bjarne Stroustrup. *The C++ Programming Language*. Addison-Wesley, 1987.

[45] Sun Microsystems. *Java Language Specification*, version 1.0 beta edition, October 1995. Available at ftp://ftp.javasoft.com/docs/javaspec.ps.zip.

[46] Tim Sweeney. The next mainstream programming language: a game developer's perspective. In *Proc. ACM Symp. on Principles of Programming Languages (POPL) '06*, page 269, January 2006.

[47] Kresten Krab Thorup. Genericity in Java with virtual types. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, number 1241 in Lecture Notes in Computer Science, pages 444–471. Springer-Verlag, 1997.

[48] Mads Torgersen. Virtual types are statically safe. In *5th Workshop on Foundations of Object-Oriented Languages (FOOL)*, January 1998.

[49] Philip Wadler. Views: A way for pattern matching to cohabit with data abstraction. In *Proc. ACM Symp. on Principles of Programming Languages (POPL) '87*, pages 307–312, January 1987.

[50] Alessandro Warth, Milan Stanojević, and Todd Millstein. Statically scoped object adaptation with expanders. In *Proc. 21st ACM Conference on Object-Oriented Programming Systems, Languages and Applications (OOP-SLA)*, Portland, OR, October 2006.

[51] Andrew K. Wright and Matthias Felleisen. A syntactic approach to type soundness. *Information and Computa-tion*, 115(1):38–94, 1994.