

Lightweight Confinement for Featherweight Java

Tian Zhao Jens Palsberg[†] Jan Vitek[†]

University of Wisconsin, Milwaukee

[†] Purdue University

ABSTRACT

Confinement properties impose a structure on object graphs which can be used to enforce encapsulation properties essential to certain program optimizations, modular reasoning, and software assurance. This paper formalizes the notion of *confined type* in the context of Featherweight Java. A static type system that mirrors the informal rules of Grothoff *et al.* [17] is proven sound. The definition of confined types is extended to confined instantiation of generic classes. This allows for confined collection types in Java and for classes that can be confined *post hoc*. Confinement type rules are given for Generic Featherweight Java, and proven sound.

1. INTRODUCTION

While object-oriented languages provide syntactic support for encapsulating fields of object structures, this form of encapsulation is only skin deep. One need only take a step back and look at the object graphs induced by the reference relationship between objects during program execution to see that they owe more to a much maligned Italian national dish than to the result of careful engineering. The problem lies in the lack of any structure on patterns of references. While access modifiers provided by most languages can prevent some part of the system from reading protected variables, there is no syntax for constraining the spread of the objects they denote.

The last five years have seen a renewed interest (starting with [8]) in the study of ownership structures as well as related work on aliasing. The goal of these works is ambitious: extend existing language with support for object encapsulation. Some promising results suggest that it is possible to capture many interesting encapsulation properties by type systems. Clarke *et al.* were the first to propose a Java-like language with ownership types [7]. More recently, Boyapati *et al.* [3, 4] have extended these results and applied to the Real-Time Specification for Java.

Our work takes a slightly different view. Rather than going for the most expressive notion of confinement, we look for the *least disruptive* one. The rationale is that before settling on one particular notion of confinement and incorporating that in a new language design, it is necessary to get first-hand experience with the benefits and costs of developing large software with these new constructs. Even more than type systems, limitations on references patterns inherent to any confinement discipline will impose constraints on pro-

grams. The risk is that they become too cumbersome for programmers without providing sufficient benefit. In [17] Bokowski and Vitek proposed a lightweight notion of encapsulation for Java called *confined types*. The idea is simple: use Java’s notion of software module (packages) as an encapsulation boundary. A class is termed *confined* if references to instances of the class may not leak out of the class’ defining package. In other words, a confined object can only be stored in fields of objects defined in the same package and manipulated by code of classes belonging to its package. What makes this approach lightweight is that very few annotations are required (one annotation per confined class, and some extra annotation for inherited method) and that conformance to the confinement rules can be checked modularly. In later work Grothoff, Palsberg and Vitek [11] designed a tool for inferring confinement without any annotations, the result of a thorough analysis of a large corpus of code (101,893 classes) is that confined classes occur naturally in large systems¹. Confinement has been shown to be applicable in practice as evidenced by the work of Clarke *et al.* on Enterprise Java Beans [6].

Two drawbacks of our previous work, highlighted by the corpus analysis, are that (a) classes can only be confined within a single package and (b) standard collection classes (such as vectors, lists) can not be used to hold confined objects. Furthermore, the informal set of rule posited in [17] was not formally shown to be correct. This paper makes the following contributions:

- Give a static type system for the rules of [17] in the context of a simple object calculus modeled on Featherweight Java [14], and prove its soundness.
- Propose a notion of generic *confined* types which addresses the two above mentioned problems.
- Give a static type system for Generic Featherweight Java with confined types, and prove its soundness.

The paper proceeds with a detailed description of confined types. Then Section 3 introduces Confined FJ (CFJ) giving it an operational semantics and a static type system. Section 4 presents our Generic Confined FJ (GCFJ). Section 5 discusses related work. Finally we conclude with some thoughts about future developments.

¹Software referred to in this paper is available from www.ovmj.org.

2. CONFINED TYPES

In object-oriented programming languages such as Java or C#, confinement can be achieved by disciplined use of built-in static access control mechanisms combined with some simple coding idioms. Confinement, as defined in [11], enforces the following informal soundness property:

An object of confined type is encapsulated in its defining scope.

To this end object types are categorized into *public types* and *confined types*. The idea is that modules are composed of two distinct software layers: an interface composed of public classes and a core consisting of confined classes (see Figure 1). What confinement adds to the visibility rules provided by the language is the guarantee that subtyping can not be used to ‘leak’ reference to core classes. Furthermore confinement annotations make the programmer’s intent explicit and allow for automated checking. In the Java current realization the unit of modularity is the package. We rely on a programming idiom to identify confined classes: marker interfaces are used to label classes and marker exceptions are used to label methods. Thus we assume the following definitions.

```
package marker;

interface ConfinedClass {}

final public class AnonymousMethod
    extends Error {}
```

Consider the following example. A class `Bucket` is used to implementation of a hash table class, `mine.HTable`. Hash table buckets are an example of internal data structures which should not escape the context of the enclosing class. In Java, the first step towards that goal is to declare class `Bucket` package scoped, thus ensuring that its visibility is restricted to the package of the `HTable` class, in this case `mine`.

```
package mine;

public class HTable {
    private Bucket[] buckets;
    public Object get(Object key) { ... }

    class Bucket implements marker.ConfinedClass {
        Bucket next;
        Object key, val; }}
```

But what if one of `HTable`’s public methods, say `get()`, were to return a `Bucket` instance or store a reference to a bucket in one of its public fields? One can view this as an escape problem: can references to instances of a package-scoped class escape their enclosing package? If the answer to this question is no, then the objects of such a class are encapsulated. This form of encapsulation is what confined types ensure. We emphasize that a stronger notion of confinement could provide even tighter bounds on reference patterns, for instance keep buckets of one hash table from being mixed up with those of another [7], but that would place a considerably greater burden on programmers.

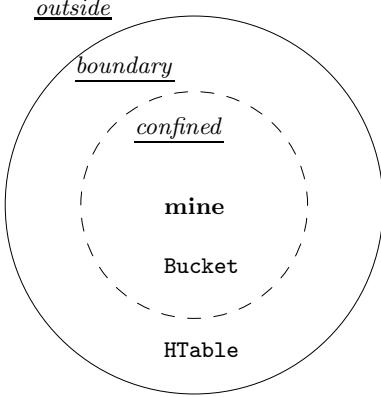


Figure 1: A package consist of a confined core and a boundary of public classes visible to the outside. Operations on a confined class can only be initiated by a classes belonging to the same package.

The pleasant characteristics of confinement is that it can be viewed as a programming discipline, *i.e.* it is possible to list a small (and simple!) set of software design rules that will result in a confined type. Moreover, these rules are local to the confining package, which implies that confined types can be used when appropriate and ignored when unnecessary.

Confinement can be enforced (or inferred) by two sets of constraints. The first set of constraints, *confinement rules*, apply to the enclosing package, the package in which the confined class is defined. These rules track values of confined types and ensure that they are neither exposed in public members, nor widened to non-confined types. The second set of constraints, so-called *anonymity rules*, applies to methods inherited by the confined classes, potentially including library code, and ensures that these methods do not leak a reference to the distinguished variable `this` which may refer to an object of confined type. There are two design choices for dealing with inherited methods either disallow inheritance or restrict it to a safe subset. We have chosen the later; anonymous method are well-behaved in the sense that they can not ‘leak’ their `this` pointer. Enforcing confinement implies tracking the spread of confined objects within a package and preventing them from crossing package boundaries. A confinement breach occur if an instance of a confined type escapes from its package. Since confinement is couched in terms of object types, widening a value from a confined type to a non-confined type presents a risk as it is not possible to ascertain what will happen with the object after the cast, thus casts are a considered to be confinement violations. We present the rules of [11] with some of implementation specific details omitted².

Anonymity Rules. Anonymity rules apply to inherited methods which may reside in classes outside of the enclos-

²We do not discuss native methods, reflection, and open packages as these are not germane to the formalization. The reader is referred to [11] for a discussion of those features and associated restrictions.

ing package. These rules prevent a method from leaking the `this` reference. A method is *anonymous* if it has the following property.

A1	The <code>this</code> reference is only used to select fields and as receiver in invocation of other anonymous methods.
----	---

This prevents an inherited method from storing or returning `this` as well as using it as an argument to a call. Selecting a field is always safe, it can not break anonymity because only the fields visible in the current class can be accessed. Method invocation (on `this`) is restricted to other methods that are anonymous as well.

Confinement Rules. The following confinement rules must hold for all classes of a package containing confined types.

C1	All inherited methods invoked on a confined type must be anonymous.
C2	A confined type cannot be public.
C3	A confined type cannot appear in the type of a public (or protected) field or the return type of a public (or protected) method of a non-confined type.
C4	Subtypes of a confined type must be confined.
C5	Confined types can only be widened to other types confined in the same package.

Rule C_1 ensures that no inherited method invoked on a confined type will leak the `this` pointer. Rule C_2 states that public classes can not be confined. Rule C_3 prevents exposure of confined types in the public interface of the package as client code could break confinement by casting them to non-confined types. Rule C_4 prevents non-confined classes (or interfaces) from extending confined types³. Finally, rule C_5 prevents values of confined type from being cast to non-confined types.

Modular enforcement. Confinement annotations can trivially be checked as part of the source level type checking or by bytecode verification. One of the key design requirements for Confined Types was that code that did not use them should not have to be checked and should not be required to use a modified compiler. Confined Types meet this requirement as the rules outlined above place no constraints on clients of a confined package (rule C_3 is crucial in this respect). Likewise confined type inference can be performed on a per-package basis. Inference of anonymous methods is somewhat trickier as it requires analyzing parent classes.

³We note that rule C_4 is not strictly needed for enforcing confinement. Rule C_5 is sufficient to ensure against leaks even if subtyping is allowed. We retain C_4 only because it will prove useful for generics.

2.1 Empirical evaluation with Kacheck/J

The Kacheck/J static analysis tool [11] was implemented to test our assumptions about confinement in real life setting. Its goal was to help understand the properties of large, well-engineered, systems. Kacheck infers confinement on unannotated Java programs by an efficient analysis of the bytecode. Since client code is not required when checking confinement, it is possible to use Kacheck on libraries. We have extended the Purdue Benchmark Suite (PBS) to 38 Java systems of varying size, purpose and origin. The entire suite contains 101,893 classes (or 282 MB of bytecode), 8,920,978 methods and 4,564 packages. To the best of our knowledge the PBS is the largest such collection of Java programs freely available (it can be obtained from www.ovmj.org). The results gleaned from the analysis are encouraging as 6,630 confined types occur in plain Java programs, Figure 2 gives the breakdown per program or library. We surmise that these numbers would be even higher if programmers actually designed code with confinement in mind. Another important statistic is that 48% of all methods (some 4,039,281) are anonymous which suggests that significant chunks of behavior can be safely inherited.

Generics. In Java programs, vectors, hashtables and other container type are omnipresent and every time an object is stored in a container, its type is widened to `Object` leading to a violation of the confinement rule preventing widening. We posit that if Java supported proper parametric polymorphism, the large majority of these violations would disappear – while truly heterogeneous data structures do exist, they seem to be the exception. If we tune our analysis to ignore widening violations linked to containers, which can be done by simply ignoring all widenings to `Object` that occur in calls to methods of classes `java.util`, 2,000 additional classes become confined. Figure 2 gives the generic confinement numbers for our new benchmark suite.

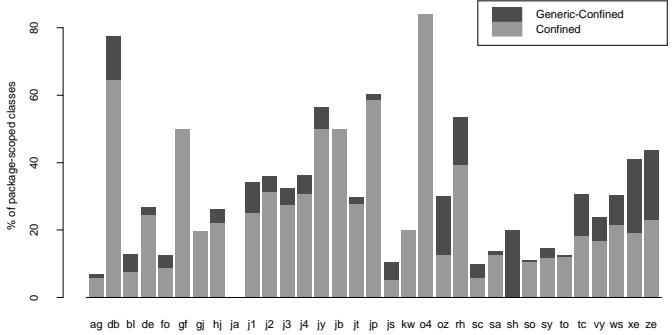


Figure 2: The Purdue Benchmark Suite. For each program, we give the ratio of confined classes and an estimate of classes that would become confined with generics.

2.2 Confinement v. Reuse

From a practical perspective, one criticism of confined types is that they seem to preclude code reuse. For a class to be confined it must be local to a particular package and, by definition, inaccessible to all other packages. This can become unwieldy when dealing with programs that require

the same logic to be available in, and confined to, different packages. While the issue was addressed in [17], a truly satisfactory remained elusive.

Any solution to the reuse problem must allow to define classes in a natural fashion, *i.e.* without imposing coding conventions more restrictive than those presented above, and must permit use of those classes as confined types in certain contexts and non-confined in other. Previous work failed to provide a satisfactory solution to this problem.

Summary. This section gave an overview of confined types. While our benchmark results have provided supporting evidence to our claim that confined types are a natural construction in object oriented systems, we are left with the following questions. Is it possible to define a variant of the confinement property that would deal with collection classes? Can we solve the reuse problem alluded to above? Moreover, until now our work has been entirely informal, can we make any soundness claims?

3. CONFINED FEATHERWEIGHT JAVA

Featherweight Java, or FJ, is core object calculus that was developed by Igarashi, Pierce and Wadler [14] for modeling Java’s type system. The calculus is minimal in that it has only five forms of expression: object creation, method invocation, field access, casting and variable access. It is in this spartan setting that Igarashi *et al.* studied a number aspects of Java, notably generics and inner classes.

Featherweight Java is a good vehicle for semantic investigations due to its economy of principles. Thus to stay true to its designer’s spirit, *Confined Featherweight Java*, or ConfinedFJ, has been modified as little as possible. In particular, we have resisted the temptation to add assignment⁴. The syntax appears on Figure 3. The only departures from FJ are the addition of package names in class declarations and of an access modifier on classes. + means the class is public and - means that the class is confined – in Java this also correspond to the default access modifier. We follow the Smalltalk-80 convention and assume that fields are package scoped by default, though adding access modifiers to fields would not be a major issue.

Consider the following ConfinedFJ program in which a public class `Table`, defined in `my` package, contains a linked list of `Bucket` objects. The `Bucket` class also belongs to `my` package but is declared confined. The program is invalid as the body of the `get()` method of `Table` returns an instance of the confined class `Bucket`, which is a direct violation of the widening rule.

⁴It may appear paradoxical that we are trying to address aliasing issues in a calculus which does not have references. But one of the insights from our previous work is that confinement inference is only dealing with the propagation of types within the program. All values of a type are treated as equal and the confinement rules ensure that values of a confined type are never confused with other values. Dealing with references is, of course, important in ownership systems where the unit of encapsulation is finer grained.

```
+ class my.Table extends Object {
    Bucket buck;
    Table(Bucket buck) {
        super(); this.buck = buck; }
    Object get() { return this.buck; } }

- class my.Bucket extends Object {
    Bucket() { super(); } }

+ class my.Factory extends Object {
    Factory() { super(); } }
    Table table() {
        return new Table( new Bucket() ); } }
```

In the above, the `Factory` class is needed to allow code outside of the `my` package to create new tables with an initial bucket (since class `Bucket` is package scoped this is the only way to use the hash table from the outside). The breach can be exhibited by constructing a class, aptly named `Breach`, in the `outside` package.

```
+ class outside.Breach extends Object {
    Object something;
    Breach(Object something) {
        super(); this.something = something; } }
```

Finally, in the context of these definitions, the expression

```
new Breach( new Factory().table().get() );
```

evaluates in two reduction steps to the following expression

```
new Breach( new Bucket() );
```

While the original expression is a valid object structure, the result isn’t. In the original, `Breach` refers to `Table` which, in turn, refers to `Bucket` enforces the desired isolation property. Namely, the confined type is only referenced by another class of the same package. In the resulting expression the confined bucket is directly held by the `Breach` object.

In another prototypical breach of confinement consider the following situation in which the confined class `Self` extends a `Broken` class which resides `outside`. Assume further that the table inherits its parent’s code for the `reveal` method.

```
- class my.Self extends outside.Broken {
    Self() { super(); } }

+ class my.Ifc extends Object {
    Ifc() { super(); }
    Object get() {
        return new Self().reveal(); } }
```

Inspection of this code does not reveal any breach of confinement. But if we widen the scope of our analysis a little and look at the `Broken` class, then we may see:

```
+ class outside.Broken extends Object {
    Broken() { super(); }
    Object reveal() { return this; } }
```

The invoking `reveal` on an instance of `Self` will return a reference to the object itself (suitably widened to `Object`).

We now proceed to introduce the syntax and semantics of ConfinedFJ, before presenting a type system which will declare both of the above programs to be ill-typed.

3.1 Syntax

The syntax of ConfinedFJ is shown in Figure 3. We use metavariables L to range over class declarations, C, D, E to range over classes, K, M to range over constructors and methods, and f and x to range over fields and parameters to methods. We use over-bar to represent a finite array, for instance, \overline{f} represents f_1, f_2, \dots, f_n . We use P, Q to range over package names and we use p_C to refer to the package that the class C is defined in. We assume a class table CT which stores the definitions of all classes and $CT(C)$ is the definition of class C .

We use e, d to range over expressions and u, v, w to range over *fully-evaluated* objects of the form `new C(̄v)`. An expression e can be either a variable x , the `this` pseudo variable, a field access $e.f$, a method invocation $e.m(\overline{e})$, a cast $(C)e$, an object `new C(\overline{e})`.

The subtyping relation $C <: D$ denotes that class C is a subtype of class D , irrespective of access mode and or package, if class C transitively extends D . We take the simplifying assumption that class names are unique.

The definitions and auxiliary functions used in our semantics and typing rules are listed in Figures 3 and 4. The predicate $visible(C, D)$ holds if the type C is *visible* in the package of D . That is, if C and D are not in the same package, then C must be declared public for it to be visible from D . This definition models the Java's class access modifiers so that only *public* classes can be referred to from other packages.

The partial order \preceq on types is such that $C \preceq D$ holds iff C is a subtype of D , and either C is public or D is confined. This definition will be used in the typing rules to prevent *reference widening*, where the reference to an object of confined type should not be widened to public types.

The *anon* predicate holds for a method m in class C , if the pseudo variable `this` is used solely for field selection and method invocation of methods that are anonymous themselves.

The other definitions are utilities mostly straight out of the FJ paper. *fields* returns the list of all fields of a class including inherited ones. *mdef* returns the name of defining class for a given method. *mtyp* returns the type signature of a method. *mbody* returns the body of the method in a given class. The predicate *override* holds if a method is valid redefinition of an inherited method.

3.2 Reduction

The dynamic semantics of our language is given by a reduction relation of the form $e \rightarrow e'$. As usual, \rightarrow^* denotes the reflexive and transitive closure of \rightarrow . The congruence rules are explicitly specified though not very surprising. There are three reductions rules. (R-FIELD) evaluates a field access expression. (R-CAST) evaluates a cast expression, and (R-INV) evaluates a method invocation expression.

ConfinedFJ has a by-value semantics which requires that arguments be fully evaluated before performing method invocations or field access. Notice that the left hand sides of the reduction rules (R-FIELD), (R-CAST), and (R-INV) contain expressions of the form `new C(̄v)`. While this is superficially closer to a real language such as Java, the reason for this choice is that it deals with values greatly simplifies the proof of the confinement property.

3.3 Typing

Typing rules are given in Figure 5, and they check for reference widening of types. In the typing rules, the condition of the form $C \preceq D$ requires that C be a subtype of D and if C is a confined type, then D must be confined as well. The typing rules for expressions are similar to those in FJ except that reference widening is not allowed. The expression typing rules emulate the effects of the confinement rules $C1$ and $C5$. In particular, Rule T-NEW prevents instantiating an object with fields of public types by arguments of confined types. Rule T-INV prevents passing arguments of confined types to a method with parameters of public types. Moreover, if the called method is defined in a class of public type while the receiver expression of the call is of confined type, then the method must be anonymous in order to prevent implicit reference widening of the variable `this`. Rule T-CAST prevents casting an expression of confined type to public type.

The definition of anonymous method is given in Figure 4. A method defined in class D is anonymous if the return expression e of the method is anonymous in D . That is, in e , the variable `this` can only be used for accessing fields and for invoking anonymous methods defined or inherited in D .

In the typing rule for methods, the return expression of a method m defined in class C must be visible in C , that is, the type of every subexpression in m must be visible in C . The effects of the constraint are that

- a method that returns expression of confined type C can not be invoked outside the package where C is defined;
- a field of confined type C can not be accessed outside the package where C is defined;
- a confined type C can not be used in cast or object instantiation outside the package where C is defined.

The visibility constraint of method body is meant to simplify our typing rules and it is not useful in practice since it requires checking the whole program. The confinement rules

Syntax:

$$\begin{aligned}
 \circ &::= + \mid - \\
 L &::= \circ \text{ class } P.C \triangleleft D \{ \bar{C} \bar{f}; K \bar{M} \} \\
 K &::= C(\bar{C} \bar{f}) \{ \text{super}(\bar{f}); \text{this.} \bar{f} = \bar{f}; \} \\
 M &::= C \bar{m}(\bar{C} \bar{x}) \{ \text{return } e; \} \\
 e &::= x \mid \text{this} \mid e.f \mid e.\bar{m}(\bar{e}) \mid (C) e \mid \text{new } C(\bar{e}) \\
 v &::= \text{new } C(\bar{v})
 \end{aligned}$$

Subtyping:

$$\begin{gathered}
 C <: C \\
 \frac{C <: D \quad D <: E}{C <: E} \\
 \frac{CT(C) = \circ \text{ class } P.C \triangleleft D \{ \dots \}}{C <: D}
 \end{gathered}$$

Computation:

$$\begin{gathered}
 \frac{\text{fields}(C_0) = (\bar{C} \bar{f}) \quad \text{new } C_0(\bar{v}).f_i \rightarrow v_i}{(\text{new } C_0(\bar{v}))} \quad (\text{R-FIELD}) \\
 \frac{C <: D}{(\text{D}) \text{ new } C(\bar{v}) \rightarrow \text{new } C(\bar{v})} \quad (\text{R-CAST}) \\
 \frac{mbody(m, C_0) = (\bar{x}, e)}{\text{new } C_0(\bar{u}).m(\bar{v}) \rightarrow [\bar{v}/\bar{x}, \text{new } C_0(\bar{u})/\text{this}]e} \quad (\text{R-INVK})
 \end{gathered}$$

Congruence:

$$\begin{gathered}
 \frac{e \rightarrow e'}{\text{new } C(\dots e \dots) \rightarrow \text{new } C(\dots e' \dots)} \\
 \frac{e \rightarrow e'}{(C) e \rightarrow (C) e'} \quad \frac{e \rightarrow e'}{e.f \rightarrow e'.f} \\
 \frac{e \rightarrow e'}{e.m(\bar{e}) \rightarrow e'.m(\bar{e})} \quad \frac{e \rightarrow e'}{e_0.m(\dots e \dots) \rightarrow e_0.m(\dots e' \dots)}
 \end{gathered}$$

Other definitions:

$$\begin{gathered}
 \frac{CT(C) = - \text{ class } P.C \triangleleft D \{ \dots \}}{\text{conf}(C)} \\
 \frac{CT(C) = + \text{ class } P.C \triangleleft D \{ \dots \}}{\text{public}(C)} \\
 \frac{\text{public}(C) \vee \rho_C = \rho_D}{\text{visible}(C, D)} \\
 \frac{C <: D \quad \text{public}(C) \vee (\text{conf}(D) \wedge \rho_C = \rho_D)}{C \preceq D}
 \end{gathered}$$

Figure 3: Confined FJ

Field look-up:

$$\begin{gathered}
 \text{fields}(\text{Object}) = () \\
 \text{fields}(D) = (\bar{D} \bar{g}) \\
 \frac{CT(C) = \circ \text{ class } P.C \triangleleft D \{ \bar{C} \bar{f}; K \bar{M} \}}{\text{fields}(C) = (\bar{D} \bar{g}, \bar{C} \bar{f})}
 \end{gathered}$$

Method definition lookup:

$$\begin{gathered}
 \frac{CT(C) = \circ \text{ class } P.C \triangleleft D \{ \bar{C} \bar{f}; K \bar{M} \} \quad B \bar{m}(\bar{B} \bar{x}) \{ \text{return } e; \} \in \bar{M}}{mdef(\bar{m}, C) = C}
 \end{gathered}$$

$$\begin{gathered}
 \frac{CT(C) = \circ \text{ class } P.C \triangleleft D \{ \bar{C} \bar{f}; K \bar{M} \} \quad \text{method } m \text{ is not defined in } \bar{M}}{mdef(m, C) = mdef(m, D)}
 \end{gathered}$$

Method type lookup:

$$\begin{gathered}
 \frac{mdef(\bar{m}, C) = D \quad CT(D) = \circ \text{ class } P.D \triangleleft E \{ \bar{C} \bar{f}; K \bar{M} \} \quad B \bar{m}(\bar{B} \bar{x}) \{ \text{return } e; \} \in \bar{M}}{mtype(\bar{m}, C) = \bar{B} \rightarrow B}
 \end{gathered}$$

Method body look-up:

$$\begin{gathered}
 \frac{mdef(\bar{m}, C) = D \quad CT(D) = \circ \text{ class } P.D \triangleleft E \{ \bar{C} \bar{f}; K \bar{M} \} \quad B \bar{m}(\bar{B} \bar{x}) \{ \text{return } e; \} \in \bar{M}}{mbody(\bar{m}, C) = (\bar{x}, e)}
 \end{gathered}$$

Valid method overriding:

$$\frac{mtype(\bar{m}, C_0) = \bar{D} \rightarrow D \Rightarrow (C = D \wedge \bar{C} = \bar{D})}{override(\bar{m}, C_0, \bar{C} \rightarrow C)}$$

Anonymous method:

$$\begin{gathered}
 \frac{mdef(\bar{m}, C) = D \quad mbody(\bar{m}, C) = (\bar{x}, e) \quad \text{anon}(e, D)}{\text{anon}(\bar{m}, C)} \\
 \frac{\text{anon}(e, C)}{\text{anon}((D) e, C)} \quad \frac{\text{anon}(\bar{e}, C)}{\text{anon}(\text{new } D(\bar{e}), C)} \quad \frac{}{\text{anon}(x, C)} \\
 \frac{}{\text{anon}(\text{this}.f, C)} \quad \frac{\text{anon}(e, C)}{\text{anon}(e.f, C)} \\
 \frac{\text{anon}(\bar{m}, C) \quad \text{anon}(\bar{e}, C)}{\text{anon}(\text{this}.m(\bar{e}), C)} \quad \frac{\text{anon}(e, C) \quad \text{anon}(\bar{e}, C)}{\text{anon}(e.m(\bar{e}), C)}
 \end{gathered}$$

Figure 4: Auxiliary Definitions

$C2$ and $C3$ achieve the same effect and they can be checked locally within the defining packages of the confined types.

The confinement rule $C4$ requires that a confined class be inherited only by the confined classes defined in the same package. We find such a constraint unnecessary in proving the main results of confinement properties. Even if a confined class C is inherited by a confined class D in another package, it is not possible for an object o of type D to be widened to the type C since our typing rules require o to be types confined in the package of D . It is possible for the public subclasses of a confined class to inherit fields of confined types or methods that return confined types, but those fields and methods are not accessible to the subclasses by Rule $C3$. Thus, confined objects can not be leaked outside their defining package through inheritance.

We assume that all classes are maintained in a class table CT and CT is well-typed if all classes in CT are well-typed. For the rest of the paper, we assume that CT is well-typed.

The two examples in the beginning of the section are not typable according to the typing rules. In the first example, the method `get` of the class `my.Table` does not type-check because Rule T-METHOD requires that if the type of the return expression of a method is confined than the return type of the method must be confined as well. The type of the return expression of the `get` method is `my.Bucket` while the return type of the method is `Object`. Since `Bucket` is confined and `Object` is public, we have a violation of the typing rule T-METHOD.

The second example does not type-check because the confined class `my.Self` inherits a non-anonymous method `reveal` from the class `outside.Broken` and the method `reveal` is called on a `my.Self` object in the `get` method of the class `my.Ifc`. The method `reveal` violates the typing rule T-INVK because the rule requires that all methods invoked on an object of confined type is either anonymous or defined in the same package as the confined type. Notice that the method `reveal` is not anonymous because the variable `this` is returned while according to the rules for anonymous methods in Figure 4, the variable `this` can only be used for field access or invocation of anonymous methods in the return expression.

3.4 Properties

In this section, we describe the properties of confined objects. Suppose we have an object o of confined type C defined in package P . The typing rules guarantee that *only methods defined in P or anonymous methods inherited by C can access the fields and methods of o*.

We prove the property in Theorem 7. We show that for a well-typed method call expression $\text{new } C_0(\bar{u}).m(\bar{v})$, where m is defined in class D_0 , all accessible objects during the evaluation of the method call are visible in the package of D_0 except that if m is anonymous, then $\text{new } C_0(\bar{u})$ is accessible for field select and invocation of anonymous methods. We also prove the usual subject reduction lemma and state the progress lemma. For the subject reduction lemma, we show

Expression typing:

$$\begin{array}{c} \frac{\Gamma \vdash x : \Gamma(x)}{\Gamma \vdash e : C_0 \quad \text{fields}(C_0) = (\bar{C} \bar{f})} \quad (\text{T-VAR}) \\ \frac{\Gamma \vdash e : C_0 \quad \Gamma \vdash \bar{e} : \bar{C} \quad \text{mtype}(m, C_0) = \bar{D} \rightarrow C \quad \bar{C} \preceq \bar{D} \quad \text{mdef}(m, C_0) = D_0 \quad C_0 \preceq D_0 \vee \text{anon}(m, D_0)}{\Gamma \vdash e.f_i : C_i} \quad (\text{T-FIELD}) \\ \frac{\Gamma \vdash e : C_0 \quad \Gamma \vdash \bar{e} : \bar{C} \quad \bar{C} \preceq \bar{D}}{\Gamma \vdash e.m(\bar{e}) : C} \quad (\text{T-INVK}) \\ \frac{\text{fields}(C) = (\bar{D} \bar{f}) \quad \Gamma \vdash \bar{e} : \bar{C} \quad \bar{C} \preceq \bar{D}}{\Gamma \vdash \text{new } C(\bar{e}) : C} \quad (\text{T-NEW}) \\ \frac{\Gamma \vdash e : D \quad \text{public}(D) \vee (\text{conf}(C) \wedge \rho_C = \rho_D)}{\Gamma \vdash (C) e : C} \quad (\text{T-CAST}) \end{array}$$

Method typing:

$$\frac{\bar{x} : \bar{C}, \text{this} : C_0 \vdash e : D \quad D \preceq C \quad \text{override}(m, D_0, \bar{C} \rightarrow C) \quad \bar{x} : \bar{C}, \text{this} : C_0 \vdash \text{visible}(e, C_0)}{C m(\bar{C} \bar{x}) \{ \text{return } e; \} \text{ OK IN } C_0 \triangleleft D_0} \quad (\text{T-METHOD})$$

Class typing:

$$\frac{\text{fields}(D) = (\bar{D} \bar{g}) \quad \bar{M} \text{ OK IN } C \triangleleft D \quad K = C(\bar{D} \bar{g}, \bar{C} \bar{f}) \{ \text{super}(\bar{g}); \text{this.}\bar{f} = \bar{f}; \} \quad \circ \text{class } P.C \triangleleft D \{ \bar{C} \bar{f}; K \bar{M} \} \text{ OK}}{} \quad (\text{T-CLASS})$$

Expression visibility:

$$\begin{array}{c} \frac{\Gamma \vdash x : C \quad \text{visible}(C, C_0)}{\Gamma \vdash \text{visible}(x, C_0)} \quad (\text{V-VAR}) \\ \frac{\Gamma \vdash \text{visible}(e, C_0) \quad \Gamma \vdash e.f_i : C \quad \text{visible}(C, C_0)}{\Gamma \vdash \text{visible}(e.f_i, C_0)} \quad (\text{V-FIELD}) \\ \frac{\Gamma \vdash e.m(\bar{e}) : C \quad \text{visible}(C, C_0) \quad \Gamma \vdash \text{visible}(\bar{e}, C_0)}{\Gamma \vdash \text{visible}(e.m(\bar{e}), C_0)} \quad (\text{V-INVK}) \\ \frac{\text{visible}(C, C_0) \quad \Gamma \vdash \text{visible}(\bar{e}, C_0)}{\Gamma \vdash \text{visible}(\text{new } C(\bar{e}), C_0)} \quad (\text{V-NEW}) \\ \frac{\text{visible}(C, C_0) \quad \Gamma \vdash \text{visible}(e, C_0)}{\Gamma \vdash \text{visible}((C) e, C_0)} \quad (\text{V-CAST}) \end{array}$$

Figure 5: Typing Rules

that an expression of public type can not be reduced to an expression of confined type. In Theorem 6, we prove that a well-type expression will not get stuck and an expression of public type will not be reduced to a value of confined type.

LEMMA 1. Subject Reduction: *If $\emptyset \vdash e : C$ and $e \rightarrow e'$ then $\emptyset \vdash e' : C'$ for some $C' \preceq C$.*

LEMMA 2. Progress: *If $\emptyset \vdash e : C$ and e contains subexpression e_0 where*

$$\begin{aligned} e_0 &= \text{new } C_0(\bar{v}).f_i \\ &\quad | \quad (C) \text{ new } C'(\bar{v}), C' <: C, \\ &\quad | \quad \text{new } C_0(\bar{u}).m(\bar{v}), \end{aligned}$$

then there exists $e' \neq e$ such that $e \rightarrow e'$.

LEMMA 3. *If $mtype(m, C_0) = \bar{D} \rightarrow D$, $mdef(m, C_0) = D_0$, and $mbody(m, C_0) = (\bar{x}, e)$, then there exists some $C \preceq D$ such that $\bar{x} : \bar{D}, \text{this} : D_0 \vdash e : C$.*

The following two lemmas prove term substitution preserves typing for expressions in anonymous methods and nonanonymous methods.

LEMMA 4. *If $\bar{x} : \bar{B}, \text{this} : D_0 \vdash e : B$, $\emptyset \vdash \bar{v} : \bar{A}$, $\bar{A} \preceq \bar{B}$, $\emptyset \vdash \text{new } C_0(\bar{u}) : C_0$, $C_0 <: D_0$, and $\text{anon}(e, D_0)$, then $\emptyset \vdash [\bar{v}/\bar{x}, \text{new } C_0(\bar{u})/\text{this}]e : C$ for some $C \preceq D$.*

LEMMA 5. *If $\bar{x} : \bar{B} \vdash e : D$, $\emptyset \vdash \bar{v} : \bar{A}$, $\bar{A} \preceq \bar{B}$. then $\emptyset \vdash [\bar{v}/\bar{x}]e : C$ for some $C \preceq D$*

In normal termination, an expression reduces to a fully-evaluated object of the form $\text{new } C(\bar{v})$. An irreducible expression of the form $(C) \text{ new } C'(\bar{v})$, where C' is not a subclass of C , represents a failed cast. An irreducible expression is *stuck* if it contains subexpressions of the form $\text{new } C_0(\bar{v}).f_i$ or $\text{new } C_0(\bar{u}).m(\bar{v})$.

THEOREM 6. *If $\emptyset \vdash e : C$ and $e \rightarrow^* e'$, then e' is not stuck, and there exists C' such that $\emptyset \vdash e' : C'$ where $C' \preceq C$.*

PROOF. Immediate from Lemma 1 and 2. \square

Intuitively, the confinement theorem states that all expressions that are manipulated within the body of a method either evaluate to methods that are visible from the method's defining package. The only exception is for anonymous methods, as they may have access to `this` which can evaluate to an instance of a class confined in another package, if this occur the method must be anonymous and the use of `this` is restricted to the receiver position.

THEOREM 7. Confinement. *If $e = \text{new } C_0(\bar{u}).m(\bar{v})$, $\emptyset \vdash e : D$, $mdef(m, C_0) = C$ and $e \rightarrow d$, then for every subexpression d' of d such that $d' \rightarrow^* \text{new } D_0(\bar{u}')$ we have:*

1. *visible(D₀, C), or*
2. *$d' = \text{new } C_0(\bar{u}')$, m is anonymous and d' appears in a subexpression of the form $d'.f$ or $d'.m'(\bar{v}')$.*

PROOF. Suppose method m is such that $mtype(m, C_0) = \bar{D} \rightarrow D$, $mdef(m, C_0) = C$, and also $mbody(m, C) = (\bar{x}, e_0)$. Let $\Gamma = \bar{x} : \bar{D}, \text{this} : C$. By Rule T-INVK, if $\emptyset \vdash \bar{v} : \bar{D}'$, then $\bar{D}' \preceq \bar{D}$; and either $C_0 \preceq C$ or $\text{anon}(m, C)$. By Lemma 3, we have that $\Gamma \vdash e_0 : D$. Let d' be any subexpression of d , there exists a subexpression e' of e_0 such that $d' = [\bar{v}/\bar{x}, \text{new } C_0(\bar{u})/\text{this}]e'$. Since $\Gamma \vdash e_0 : D$, there exists D' such that $\Gamma \vdash e' : D'$.

By Rule T-METHOD, we have $\Gamma \vdash \text{visible}(e_0, C)$ and by simple induction, we can show that $\Gamma \vdash \text{visible}(e', C)$ which implies $\text{visible}(D', C)$.

If $C_0 \preceq C$, then from $\bar{D}' \preceq \bar{D}$ and Lemma 5, there exists D'' such that $\emptyset \vdash d' : D''$ and $D'' \preceq D'$. By Lemma 1, if $d' \rightarrow^* \text{new } D_0(\bar{u}')$, then $D_0 \preceq D''$. Together with $D'' \preceq D'$ and $\text{visible}(D', C)$, we have $\text{visible}(D_0, C)$.

If we have $\text{anon}(m, C)$, then by definition, we have $\text{anon}(e_0, C)$ as well. By the definition of anonymous expression, either we have $\text{anon}(e', C)$ or $e' = \text{this}$ appears in a subexpression of the form $\text{this}.f$ or $\text{this}.m'(\bar{v}')$. If $\text{anon}(e', C)$, then from $C_0 <: C$, $\bar{D}' \preceq \bar{D}$, and Lemma 4, there exists D'' such that $\emptyset \vdash d' : D''$ and $D'' \preceq D'$. By Lemma 1, if $d' \rightarrow^* \text{new } D_0(\bar{u}')$, then $D_0 \preceq D''$. Similarly, we have $\text{visible}(D_0, C)$.

Therefore, if $d' \rightarrow^* \text{new } D_0(\bar{u}')$ then either $\text{visible}(D_0, C)$ or $d' = \text{new } C_0(\bar{u}')$ appears in a subexpression of the form $d'.f$ or $d'.m'(\bar{v}')$. \square

This result ensures that any well-typed program consisting of a top-level method invocation expression will preserve confinement.

4. GENERICS AND CONFINEMENT

The lack of support for collections and the reuse problem were identified early on as significant issues limiting the applicability of confined types. In this section we show how to extend the confinement property to generic types. In particular, how to write a generic class $C\langle X \rangle$ which can be used to define generic collection classes that are, in and of themselves, not confined. But that become confined if instantiated with confined arguments.

Following the lead of Featherweight Generic Java, ConfinedFGJ is extended with support for generic types and renamed ConfinedFGJ. The only departure from our role model is that in ConfinedFGJ, type variables are annotated with *confinability* tags. So for some generic type $C\langle X \rangle$, the type variable X can be either *public* X^P or *confined* X^C . The semantics of these annotations is that a public type variable can only be

instantiated with a public type, while a confined type variable *may* be instantiated with a confined class (or a public class). Besides the first five rules already presented, we require:

<i>C6</i>	A generic type or type variable can not be widened to a type containing a different set of confined type variables.
<i>C7</i>	A confined type can not replace a public type variable in the instantiation of a generic type.
<i>C8</i>	Overriding must preserve anonymity of methods.

C5 and *C6* combined corresponds to the subtyping partial order that prevents reference widening for Generic ConfinedFJ. *C7* corresponds to the extra requirement in the definition of well-formed generic types. Unlike in ConfinedFJ, *C8* is necessary since we are not certain which method may be called before a generic class is instantiated.

Consider a generic linked list class. If we desire to allow the class to be used to hold confined objects, it should be defined with a confined type variable in place of its value type. The following is likely definition.

```
+ class some.List<Xc extends Object>
  extends Object {
  X val;
  List<X> next;
  List(X val, List next) {
    super(); this.val=val; this.next=next;
  }
}
```

Where things get interesting is that we can now write code that uses lists in several contexts. Thus it is possible to use the list twice within the same package, once with a confined type – thus turning the list itself into a confined type – and once with a non confined type. The following example illustrates this. Classes **A** and **B** reside in **my** package, the latter is confined. Class **A** further defines two variables: **show** holds a list of **A** objects and **hide** holds a list of **B** objects. Since **B** is confined the type **List** will be confined as well.

```
+ class my.A extends Object {
  List<A> show;
  List<B> hide;
  ...
}

- class my.B extends Object {
  B() { super(); }
}
```

Generics can be used to address the reuse problem with a slight of hand. If a class needs to be reused across different packages and confined in each of these packages one may simply give the class a dummy type variable annotated as confined. This type variable need not be used in the body of

the class, it will merely serve as a marker. Reuse is obtained by instantiating the class in each of the packages with a dummy confined class as argument. Consider the following scenario, a class **Key** is meant to provide functionality that can be used in different confined settings.

```
+ class a.Key<Xc extends Object> extends Object {
  ...
}
```

The type variable **X** need not be used in the implementation of **Key**. The class can be confined in any package as long as it is instantiated with a confined type, e.g. **new Key<my.B>()**.

Generic types that may be instantiated with a generic argument are referred to as *confinable* generic types. The semantics of confinable generic types are surprisingly simple. Any type variable that is tagged as confined will be treated as a confined type by the type system and reference widening will be forbidden for expressions of this type. Even though a confinable generic type may not be confined in any package, reference widening should not be allowed for expressions of the type either.

For example, consider a generic container class.

```
+ class some.Container<Xc extends Object>
  extends Object {
  X val;
  Container(X val) {this.val = val}
  Object get() { return this.val; }
  Object get2() { return this; }
}
```

The **Container** class has a method **get** that returns the field **val** and a method **get2** that returns the pointer **this**. Both methods violate the confinement properties because the types of the return expressions in **get** and **get2** are widened from **X** and **Container<X>** to **Object** respectively. Since **X** is tagged as confined type variables, it may be replaced by confined types when **Container** is instantiated as in the class **my.A**.

```
+ class my.A extends Object {
  Container<B> f = new Container<B>(new B());
  Object reveal() {
    return f.get();
  }
  Object reveal2() {
    return f.get2();
  }
}
```

The class **my.A** has a field **f** that points to a **Container** object instantiated with a confined object of the type **my.B**. The method **reveal** calls the **get** method on the **Container** object. Apparently, the type of the expression **new my.B()** is widened from **my.B** to the type **Object** by the method **reveal** because of the call to the **get** method. The method **reveal2** also violates confinement properties because it widens confined object of the type **Container<my.B>** to **Object**. Thus,

our typing rules disallow methods such as `get` and `get2` in classes of confinable generic types to prevent widening of references to confined objects.

Any confinable generic type may be confined in more than one packages. For instance, if a class $P.A(X)$ contains a reference to an object of the type $Q.B(P.C, X)$, where $P.C$ is a confined type in P , then $Q.B(P.C, X)$ should be confined in P . If the class $P.A(X)$ is instantiated with a confined type $P'.C'$, then an object of the type $P.A(P'.C')$ is confined in P' and it also contains a reference to an object of the type $Q.B(P.C, P'.C')$, which is now confined in both P and P' .

In the previous section we argued that the rule limiting subtyping of confined types was redundant. This is not the case when generics are added to ConfinedFJ. $C4$ is needed since a generic class may contain fields of confined types. However, we can not force all fields of confined variable types to be declared private or default as this would make such fields inaccessible when the generic class is instantiated in other packages with non-confined arguments types. This is undesirable as it limits the ability to reuse generic classes.

4.1 Syntax

The syntax for generic class is shown in Figure 6. For simplicity, we omit the parameterized types in methods. We use X, Y to represent type variables and we use N, W to represent types other than type variables; We use S, T to range over types and type variables. The type variable X in a generic class can be either *public* X^P or *confined* X^c . In a class definition of the form $\circ \text{class } P.C(\bar{X} \triangleleft \bar{N}) \triangleleft N \{ \dots \}$, \bar{N} are the upper bounds for the type variables \bar{X}

Syntax:

- $\circ ::= + \mid -, N ::= C(\bar{T}), T ::= X \mid N, X ::= X^c \mid X^P$
- $L ::= \circ \text{class } P.C(\bar{X} \triangleleft \bar{N}) \triangleleft N \{ \bar{T} f; K \bar{M} \}$
- $K ::= C(\bar{T} \bar{f}) \{ \text{super}(\bar{f}); \text{this.}\bar{f} = \bar{f}; \}$
- $M ::= T m(\bar{T} \bar{x}) \{ \text{return } e; \}$
- $e ::= x \mid \text{this} \mid e.f \mid e.m(\bar{e}) \mid (N) e \mid \text{new } N(\bar{e})$
- $v ::= \text{new } N(\bar{v})$

Figure 6: Syntax

4.2 Typing

Figure 7 contains subtyping rules, definitions for well-formed types, and other miscellaneous definitions. The subtyping rules are the same as those in Generic FJ.

A type $C(\bar{T})$ is public iff C is public and \bar{T} are public types or public type variables. We use $\text{conf}(N, P)$ to denote that the type N is confined in the package P . A type $C(\bar{T})$ is confined in P iff C is confined in P or there exists a type $N \in \bar{T}$ such that $\text{conf}(N, P)$ is true. Note that a confinable generic type may not be be confined in any package. However, it can become confined in some packages after instantiation. Thus, it is necessary to prevent reference widening for expressions of confinable generic types. A type N is confined iff there exists P such that N is confined in P .

Subtyping:

$$\frac{\Delta \vdash T <: U}{\Delta \vdash S <: T \quad \Delta \vdash T <: U} \frac{}{\Delta \vdash S <: U}$$

$$\Delta \vdash X <: \Delta(X)$$

$$\frac{CT(C) = \circ \text{class } P.C(\bar{X} \triangleleft \bar{N}) \triangleleft N \{ \dots \}}{\Delta \vdash C(\bar{T}) <: [\bar{T}/\bar{X}]N}$$

Public, Confined Types and Type Variables:

$$\frac{\text{conf}(X^c) \quad \text{public}(X^P)}{\text{conf}(C(\bar{T}), P)}$$

$$\frac{(\text{conf}(C) \wedge \rho_C = P) \vee \exists T \in \bar{T}, \text{conf}(T, P)}{\text{conf}(C(\bar{T}), P)}$$

$$\frac{\text{conf}(C(\bar{T}), P) \text{ for some } P}{\text{conf}(C(\bar{T}))}$$

$$\frac{\text{public}(C) \wedge \forall T \in \bar{T}, \text{public}(T)}{\text{public}(C(\bar{T}))}$$

Well-formed Types:

$$\frac{\Delta \vdash \text{Object}}{\Delta \vdash X}$$

$$\frac{X \in \text{dom}(\Delta)}{\Delta \vdash X}$$

$$\frac{CT(C) = \circ \text{class } P.C(\bar{X} \triangleleft \bar{N}) \triangleleft N \{ \dots \}}{\Delta \vdash C(\bar{T})}$$

$$\frac{\Delta \vdash \bar{T} \quad \Delta \vdash \bar{T} <: [\bar{T}/\bar{X}]\bar{N} \quad \forall i, \text{public}(T_i) \vee \text{conf}(X_i)}{\Delta \vdash C(\bar{T})}$$

Set of Confined Type Variables in Types:

$$\text{ConfVar}(X) = \{X \mid \text{conf}(X)\}$$

$$\text{ConfVar}(C(\bar{T})) = \bigcup_{T \in \bar{T}} \text{ConfVar}(T)$$

Subtyping without Reference Widening:

$$\frac{\Delta \vdash S <: T \quad \text{conf}(S, P) \Rightarrow \text{conf}(T, P)}{\Delta \vdash S \preceq T}$$

Visibility of Types and Type Variables

$$\frac{\text{conf}(N, P) \Rightarrow P = \rho_C}{\text{visible}(N, C)}$$

$$\frac{\text{conf}(N, P) \Rightarrow (P = \rho_C \vee \text{conf}(C(\bar{N}), P))}{\text{visible}(N, C(\bar{N}))}$$

Figure 7: Subtyping rules, Well-formed types, and Miscellaneous Definitions

If the type $C(\bar{T})$ is well-formed, then for all i , either T_i is public or the type variable it replaces must be confined. This restriction prevents reference widening when generic types are instantiated.

As in confined FJ, we define a partial order \preceq on types

Bound of type:

$$bound_{\Delta}(x) = \Delta(x) \quad bound_{\Delta}(N) = N$$

Field look-up:

$$fields(\text{Object}) = ()$$

$$\frac{CT(C) = \circ \text{ class } P.C\langle \bar{X} \triangleleft \bar{N} \rangle \triangleleft N \{ \bar{S} \bar{f}; K \bar{M} \} \\ fields([\bar{T}/\bar{X}]N) = (\bar{U} \bar{g})}{fields(C\langle \bar{T} \rangle) = (\bar{U} \bar{g}, [\bar{T}/\bar{X}] \bar{S} \bar{f})}$$

Method Definition Lookup:

$$\frac{CT(C) = \circ \text{ class } P.C\langle \bar{X} \triangleleft \bar{N} \rangle \triangleleft N \{ \bar{S} \bar{f}; K \bar{M} \} \\ m \text{ is defined in } \bar{M}}{mdef(m, C\langle \bar{T} \rangle) = [\bar{T}/\bar{X}]N}$$

$$\frac{CT(C) = \circ \text{ class } P.C\langle \bar{X} \triangleleft \bar{N} \rangle \triangleleft N \{ \bar{S} \bar{f}; K \bar{M} \} \\ m \text{ is not defined in } \bar{M}}{mdef(m, C\langle \bar{T} \rangle) = mdef(m, [\bar{T}/\bar{X}]N)}$$

Method Type Lookup:

$$\frac{mdef(m, N) = D\langle \bar{T} \rangle \\ CT(D) = \circ \text{ class } P.D\langle \bar{X} \triangleleft \bar{N} \rangle \triangleleft N \{ \bar{S} \bar{f}; K \bar{M} \} \\ U \in \bar{M}}{mtype(m, N) = [\bar{T}/\bar{X}]U \rightarrow [\bar{T}/\bar{X}]U}$$

Method body look-up:

$$\frac{mdef(m, N) = D\langle \bar{T} \rangle \\ CT(D) = \circ \text{ class } P.D\langle \bar{X} \triangleleft \bar{N} \rangle \triangleleft N \{ \bar{S} \bar{f}; K \bar{M} \} \\ U \in \bar{M}}{mbody(m, N) = (\bar{x}, [\bar{T}/\bar{X}]e)}$$

Valid method overriding

$$\frac{mtype(m, N_0) = \bar{U} \rightarrow U \wedge T = U, \bar{T} = \bar{U}}{override(m, N_0, \bar{T} \rightarrow T)}$$

Figure 8: Auxiliary functions

to represent subtyping relation without reference widening. To have $\Delta \vdash S \preceq T$, we must have $\Delta \vdash S <: T$ and if S is confined in package P then so is T ; also, S, T must contain the same set of confined type variables. With the last restriction, the partial order on S, T still holds even if type variables in S, T are replaced by confined types. We use $ConfVar(N)$ to denote the set of confined type variables in N .

A type variable is visible in any class, while a non-variable type N is visible in the class C iff $conf(N, P)$ implies that C is defined in P . A type N is visible in $C\langle \bar{N} \rangle$ iff $conf(N, P)$ implies either C is defined in P or $conf(C\langle \bar{N} \rangle, P)$.

Figure 8 contains the helper functions used in the typing rules and they are similar to those in Generic FJ.

Anonymous methods of a generic class $C\langle \bar{X} \rangle$ stay anonymous even if the type parameters \bar{X} in class C are replaced by some types. In the rest of the paper, $anon(m, C\langle \bar{T} \rangle)$ is equivalent to $anon(m, C)$.

Expression typing:

$$\Delta; \Gamma \vdash x : \Gamma(x) \quad (\text{GT-VAR})$$

$$\frac{\Delta; \Gamma \vdash e_0 : T_0 \\ fields(bound_{\Delta}(T_0)) = (\bar{T} \bar{f})}{\Delta; \Gamma \vdash e_0.f_i : T_i} \quad (\text{GT-FIELD})$$

$$\frac{\Delta; \Gamma \vdash e_0 : T_0 \quad \Delta; \Gamma \vdash \bar{e} : \bar{V} \\ mdef(m, bound_{\Delta}(T_0)) = N_0 \\ mtype(m, N_0) = \bar{U} \rightarrow U \quad \Delta \vdash \bar{V} \preceq \bar{U} \\ \Delta \vdash T_0 \preceq N_0 \vee anon(m, N_0)}{\Delta; \Gamma \vdash e_0.m(\bar{e}) : U} \quad (\text{GT-INVK})$$

$$\frac{\Delta \vdash N \quad fields(N) = (\bar{T} \bar{f}) \\ \Delta; \Gamma \vdash \bar{e} : \bar{S} \quad \Delta \vdash \bar{S} \preceq \bar{T}}{\Delta; \Gamma \vdash \text{new } N(\bar{e}) : N} \quad (\text{GT-NEW})$$

$$\frac{\Delta; \Gamma \vdash e : T \quad \Delta \vdash N \\ conf(T, P) \Rightarrow conf(N, P) \\ ConfVar(T) = ConfVar(N)}{\Delta; \Gamma \vdash (N)e : N} \quad (\text{GT-CAST})$$

Method typing:

$$\frac{\Delta = \bar{X} <: \bar{N} \quad \Gamma = \bar{x} : \bar{T}, \text{this} : C\langle \bar{X} \rangle \\ \Delta \vdash \bar{T}, T \quad \Delta; \Gamma \vdash e : S \quad \Delta \vdash S \preceq T \\ \Delta; \Gamma \vdash \text{visible}(e, C) \quad override(m, N, \bar{T} \rightarrow T) \\ anon(m, N) \Rightarrow anon(m, C\langle \bar{X} \rangle) \\ T \in \bar{M}}{T \in \bar{M} \quad \Delta; \Gamma \vdash e : T} \quad (\text{GT-METHOD})$$

Class typing:

$$\frac{\bar{X} <: \bar{N} \vdash \bar{N}, N, \bar{T} \quad \bar{M} \text{ OK IN } C\langle \bar{X} \triangleleft \bar{N} \rangle \triangleleft N \\ fields(N) = (\bar{U} \bar{g}) \\ K = C(\bar{U} \bar{g}, \bar{T} \bar{f}) \{ \text{super}(\bar{g}); \text{this}.\bar{f} = \bar{f}; \} \\ \circ \text{ class } P.C\langle \bar{X} \triangleleft \bar{N} \rangle \triangleleft N \{ \bar{T} \bar{f}; K \bar{M} \} \text{ OK}}{\circ \text{ class } P.C\langle \bar{X} \triangleleft \bar{N} \rangle \triangleleft N \{ \bar{T} \bar{f}; K \bar{M} \} \text{ OK}} \quad (\text{GT-CLASS})$$

Expression visibility:

$$\frac{\Delta; \Gamma \vdash x : T \quad \text{visible}(T, C)}{\Delta; \Gamma \vdash \text{visible}(x, C)} \quad (\text{GV-VAR})$$

$$\frac{\Delta; \Gamma \vdash \text{visible}(e, C) \\ \Delta; \Gamma \vdash e.f_i : T \quad \text{visible}(T, C)}{\Delta; \Gamma \vdash \text{visible}(e.f_i, C)} \quad (\text{GV-FIELD})$$

$$\frac{\Delta; \Gamma \vdash e.m(\bar{e}) : T \quad \text{visible}(T, C) \\ \Delta; \Gamma \vdash \text{visible}(e, C) \quad \Delta; \Gamma \vdash \text{visible}(\bar{e}, C)}{\Delta; \Gamma \vdash \text{visible}(e.m(\bar{e}), C)} \quad (\text{GV-INVK})$$

$$\frac{\text{visible}(N, C) \quad \Delta; \Gamma \vdash \text{visible}(\bar{e}, C)}{\Delta; \Gamma \vdash \text{visible}(\text{new } N(\bar{e}), C)} \quad (\text{GV-NEW})$$

$$\frac{\text{visible}(N, C) \quad \Delta; \Gamma \vdash \text{visible}(e, C)}{\Delta; \Gamma \vdash \text{visible}((N)e, C)} \quad (\text{GV-CAST})$$

Figure 9: Typing rules

Figure 9 contains typing rules for expressions, methods, and classes, and also visibility rules for expressions. The expression typing rules are similar to those in Generic FJ with some additional constraints to prevent explicit or implicit reference widening.

The constraint corresponds to the confinement rule $C4$ is not needed in proving the confinement properties for generic types. However, we do need the rule if we want modular checking.

In the method typing rule, we require that the return expression of a method in class C be visible in C . The visibility rules of expressions in a generic class are similar to those of non-generic class. As in the case of CFJ, the visibility constraint of method body is meant to simplify the typing rules and it requires whole program analysis. To enable modular checking, we can substitute visibility constraint of method body with the confinement rules $C2$, $C3$, and $C4$.

Unlike in CFJ, Rule $C4$ is needed because Rules $C2$, $C3$ do not apply to fields of variable types or methods with variable return types.

For example, consider a generic confinable class $\text{my}.A<X^c>$ and its instantiation $\text{my}.A<C>$. $\text{my}.A<C>$ is confined because the type argument $\text{my}.C$ is confined. The class $\text{my}.B$ extends $\text{my}.A<C>$ and the inherited field f in the class $\text{my}.B$ is assigned a confined object `new C()` by the method m in class $\text{my}.D$. The method `breach` in the class `outside.E` is able to access the field f of `new B()` and return the confined object `new C()`. Since the class `outside.E` is defined outside the package of $\text{my}.C$, we have a violation of confinement property.

```
+ class my.A<X^c> {
    X f;
    A(X f) {this.f = f;}
}
- class my.C {}
+ class my.B extends A<C> {}
+ class my.D {
    B m() { return new B(new C()); }
}
+ class outside.E {
    Object breach(){return ((new D()).m()).f;}
}
```

The method `breach` is not typable because the Rule (T-METHOD) requires the type of `((new D()).m()).f` to be visible in `outside.E`, which is not true.

Rules $C2$, $C3$ alone are not able to prevent this problem because the type of the field f is variable type X . We may not want to set the access modifier of f as private or default, because the class $\text{my}.A<X^c>$ would be less reusable because f becomes inaccessible when $\text{my}.A<X^c>$ is instantiated outside the package my . However, according to Rule $C4$, the class $\text{my}.B$ should be confined because $C4$ requires the subclasses of a confined class to be confined. If $\text{my}.B$ were confined,

then `outside.E` could not call the method m of the class $\text{my}.D$ since m would not be accessible to `outside.E` according to Rule $C3$.

Also in the typing rule for methods, we require that method-overriding preserves anonymity for generic classes. This is needed for preventing implicit reference widening in a generic class after it is instantiated.

4.3 Properties

In this section, we prove some results similar to those for Confined FJ (CFJ). In Generic CFJ, a program without free type variables should have the same confinement properties as does a program in CFJ.

Our typing rules for generic classes guarantee that if the type of an expression e is confined in a package P and $e \rightarrow^* e'$, then the type of e' is confined in P as well. This property is shown in the subject reduction lemma and theorem 10.

LEMMA 8. Subject Reduction: *If $\emptyset; \emptyset \vdash e : N$ and $e \rightarrow e'$ then $\exists N'$ such that $\emptyset; \emptyset \vdash e' : N'$, where $\emptyset \vdash N' \preceq N$.*

PROOF. We prove by induction over the structure of e . The proof is similar to that of Lemma 1 except for the case when e is of the form `new N0(\bar{u}).m(\bar{v})`. \square

LEMMA 9. Progress: *If $\emptyset; \emptyset \vdash e : N$ and e contains subexpression e_0 where*

$$\begin{aligned} e_0 &= \text{new } N_0(\bar{v}).f_i \\ &\quad | \quad (N) \text{ new } N'(\bar{v}), \emptyset \vdash N' <: N, \\ &\quad | \quad \text{new } N_0(\bar{u}).m(\bar{v}), \end{aligned}$$

then there exists $e' \neq e$ such that $e \rightarrow e'$.

THEOREM 10. *If $\emptyset; \emptyset \vdash e : N$ and $e \rightarrow^* e'$, then e' is not stuck and $\exists N'$ such that $\emptyset; \emptyset \vdash e' : N'$, where $\emptyset \vdash N' \preceq N$.*

PROOF. Immediate from Lemma 8 and 9. \square

Theorem 11 is similar to Theorem 7. The difference is that if a method defined in the class $C\langle \bar{N} \rangle$ can access an object confined in the package P , then either C is defined in P or there exists $N \in \bar{N}$ such that N is confined in P . The latter case may be true when a generic type $C\langle \bar{X} \rangle$ is instantiated with confined type arguments.

THEOREM 11. *If $e = \text{new } N_0(\bar{u}).m(\bar{v})$, $\emptyset; \emptyset \vdash e : N$, $mdef(m, N_0) = N'_0$, and $e \rightarrow d$, then for every subexpression d' of d such that $d' \rightarrow^* \text{new } N'(\bar{u}')$ we have:*

1. *either $\text{visible}(N', N'_0)$, or*
2. *$d' = \text{new } N_0(\bar{u})$, m is anonymous, and d' appears in a subexpression of the form $d'.f$ or $d'.m'(\bar{v}')$.*

PROOF. Similar to the proof of Theorem 7. \square

5. RELATED WORK

Reference semantics permeate object-oriented programming languages, and the issue of controlling aliasing has been the focus of numerous papers in the recent years [13, 12, 9, 1, 16, 10, 15, 8]. In [16], flexible alias protection is presented as a means to control potential aliasing amongst components of an aggregate object. Clarke, Potter, and Noble [8] have formalized representation containment by means of ownership types. Boyland, Noble and Retert [5] introduced capabilities as a uniform system to describe restrictions imposed on references. Recent work by Boyapati *et al.* [3, 4] has applied ownership types to scoped memory in the Real-time Specification for Java. The paper by Banerjee and Naumann [2] demonstrated the use of object confinement to achieve representation independence. Their stronger notation of confinement is instance-based and it can be used to prove equivalence of class implementations such that if an implementation is confined, then it may be replaced by semantically equivalent ones without affecting the behavior of the whole program. Their work has significance in proving the equivalence of programs and the correctness of static analysis such as secure information flow.

In related work Clarke *et al.* [6] have adapted the notion of confined types to ensuring the integrity of components in the Enterprise JavaBeans framework. One of the interesting aspects of their work is that the unit of confinement is different. Rather than confining types within a package, they confine them within a Bean using the following rules (CB1-6): CB1 declares which types are confined (*C2* in our case), CB2 prevents confined types from appearing at the Bean boundary or in static variables (roughly equivalent to *C3*), CB3 prevents widening of confined types (identical to *C2*), CB4 prevents unconfined types to be cast to confined types (redundant as it follows from CB3 and CB6 that no cast from an unconfined type to a confined will ever succeed), CB5 prevents confined code from accessing unconfined classes which have confined types in their signature, and finally CB6 states that confined classes may only extend one another or `Object` (a stronger version of *C4*). Rule CB6 precludes confined classes from inheriting code from non-confined classes and thus sidesteps the issue of anonymous methods. But the main difference between the systems lies in rules CB2 and CB5 which conspire to prevent the use of static variables to communicate across beans.

6. CONCLUSION

This paper has formalized the notion of confined type in the context of an minimal object calculus modeled after Featherweight Java. A static type system that mirrors the informal rules confinement was proposed and proven sound. The confinement invariant was shown to hold for well-typed programs. In the second part of the paper, definition of confined types was extended to confined instantiation of generic classes. This allows for confined collection types in Java and for classes that can be confined *post hoc*. Confinement type rules are given for Generic Featherweight Java, and proven sound. A generic confinement invariant is established and proven for well-typed programs. In future work we plan to study the relationship between confined types and general purpose ownership type systems.

Acknowledgments. The authors thank James Noble for frequent inspiration and encouragements, Dave Clarke for one fewer reduction rule, Christian Grothoff for his boundless energy in generating experimental data, John Boyland for explaining borrowing, and Phil Wadler for the initial impetus to look at generics. Finally we thank the anonymous referees and attendees of IWACO for their comments.

7. REFERENCES

- [1] Paulo Sérgio Almeida. Balloon types: Controlling sharing of state in data types. In Mehmet Aksit and Satoshi Matsuoka, editors, *ECCOP'97—Object-Oriented Programming, 11th European Conference*, volume 1241 of *LNCS*, pages 32–59, Jyväskylä, Finland, 9–13 June 1997. Springer-Verlag.
- [2] Anindya Banerjee and David A. Naumann. Representation independence, confinement and access control [extended abstract]. *ACM SIGPLAN Notices*, 37(1):166–177, January 2002.
- [3] Chandrasekhar Boyapati, Robert Lee, and Martin Rinard. Ownership types for safe programming: Preventing data races and deadlocks. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, November 2002.
- [4] Chandrasekhar Boyapati, Alexandru Salcianu, William Beebee, and Martin Rinard. Ownership types for safe region-based memory management in real-time java. June 2003.
- [5] John Boyland, James Noble, and William Retert. Capabilities for aliasing: A generalisation of uniqueness and read-only. In *ECCOP'01 — Object-Oriented Programming, 15th European Conference*, number 2072 in Lecture Notes in Computer Science, Berlin, Heidelberg, New York, 2001. Springer.
- [6] Dave Clarke, Michael Richmond, and James Noble. Saving the world from bad Beans: Deployment-time confinement checking. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, Anaheim, CA, November 2003.
- [7] David Clarke. *Object Ownership and Containment*. PhD thesis, School of Computer Science and Engineering, University of New South Wales, Sydney, Australia, 2001.
- [8] David G. Clarke, John M. Potter, and James Noble. Ownership types for flexible alias protection. In *OOPSLA '98 Conference Proceedings*, volume 33(10) of *ACM SIGPLAN Notices*, pages 48–64. ACM, October 1998.
- [9] D. Detlefs, K. Rustan M. Leino, and G. Nelson. Wrestling with rep exposure. Technical report, Digital Equipment Corporation Systems Research Center, 1996.

- [10] Daniela Genius, Martin Trapp, and Wolf Zimmermann. An approach to improve locality using Sandwich Types. In *Proceedings of the 2nd Types in Compilation workshop*, volume LNCS 1473, Kyoto, Japan, March 1998. Springer Verlag.
- [11] Christian Grothoff, Jens Palsberg, and Jan Vitek. Encapsulating objects with confined types. In *Proceedings of the ACM Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA)*, 2001.
- [12] John Hogg. Islands: Aliasing Protection in Object-Oriented Languages. In *Proceedings of the OOPSLA '91 Conference on Object-Oriented Programming Systems, Languages and Applications*, pages 271–285, November 1991. Published as ACM SIGPLAN Notices, volume 26, number 11.
- [13] John Hogg, Doug Lea, Alan Wills, Dennis de Champeaux, and Richard Holt. The Geneva convention on the treatment of object aliasing. *OOPS Messenger*, 3(2), April 1992.
- [14] Atsushi Igarashi, Benjamin C. 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, May 2001.
- [15] S.J.H. Kent and I. Maung. Encapsulation and Aggregation. In *Proceedings of TOOLS PACIFIC 95 (TOOLS 18)*. Prentice Hall, 1995.
- [16] James Noble, Jan Vitek, and John Potter. Flexible alias protection. In Eric Jul, editor, *ECOOP'98—Object-Oriented Programming*, volume 1445 of *Lecture Notes In Computer Science*, pages 158–185, Berlin, Heidelberg, New York, July 1998. Springer-Verlag.
- [17] J. Vitek and B. Bokowski. Confined types in Java. *Software Practice and Experience*, 31(6):507–532, 2001.