

Expressive Modular Linking for Object-Oriented Languages

UUCS-02-014

Sean McDirmid, Wilson C. Hsieh, Matthew Flatt
School of Computing
University of Utah
{mcdirmid,mflatt,wilson}@cs.utah.edu

October 11, 2002

In this paper we show how modular linking of program fragments can be added to statically typed, object-oriented (OO) languages. Programs are being assembled out of separately developed software components deployed in binary form. Unfortunately, mainstream OO languages (such as Java) still do not provide support for true modular linking. Modular linking means that program fragments can be separately compiled and type checked, and that linking can ensure global program type correctness without analyzing program fragment implementations.

Supporting modular linking in OO languages is complicated by two expressive features that current OO languages do not support together: mixin-style inheritance across program fragment boundaries, and cyclic dependencies between program fragments. In a previous paper at OOPSLA 2001, we have demonstrated the practical uses for such expressiveness. When such expressiveness is permitted, link-time type checking rules must ensure that method collisions and inheritance cycles do not occur after program fragments are linked into a program. In this paper, we show how modular linking with both cyclic linking and mixin-style inheritance can be supported using a type-checking architecture that can be added on top of existing OO languages, such as Java.

Expressive Modular Linking for Object-Oriented Languages

Sean McDirmid, Wilson C. Hsieh, Matthew Flatt

Abstract

In this paper we show how modular linking of program fragments can be added to statically typed, object-oriented (OO) languages. Programs are being assembled out of separately developed software components deployed in binary form. Unfortunately, mainstream OO languages (such as Java) still do not provide support for true modular linking. Modular linking means that program fragments can be separately compiled and type checked, and that linking can ensure global program type correctness without analyzing program fragment implementations.

Supporting modular linking in OO languages is complicated by two expressive features that current OO languages do not support together: mixin-style inheritance across program fragment boundaries, and cyclic dependencies between program fragments. In a previous paper at OOPSLA 2001, we have demonstrated the practical uses for such expressiveness. When such expressiveness is permitted, link-time type checking rules must ensure that method collisions and inheritance cycles do not occur after program fragments are linked into a program. In this paper, we show how modular linking with both cyclic linking and mixin-style inheritance can be supported using a type-checking architecture that can be added on top of existing OO languages, such as Java.

1 Introduction

Linking is the process of turning a collection of program fragments into a complete program. **Modular** linking is important, because programs are being assembled out of separately developed components deployed in binary form. Modular linking requires that the implementation of a program fragment (fragment, for short) be compiled and type-checked separately from foreign fragments [4]. Modular linking has many benefits: it simplifies dependency management, makes program compilation and linking more efficient, promotes local reasoning of fragments, makes dynamic linking of fragments more straightforward, and makes compilation and linking errors easier for users to reason about because the errors are not related to hidden implementation details. Modular linking has been extensively studied in the context of functional languages, where function constructs especially facilitate modular linking, but modular linking for object-oriented (OO) languages is still an active area of research.

Classes in OO languages encapsulate and permit the reuse of object implementations using instantiation and inheritance. Classes and fragments address different development concerns and so should be supported with distinct constructs [21], where a fragment can conceptually be a set of classes. Unlike fragments that contain functions, modular linking of fragments that contain classes is difficult: class inheritance can create complex dependencies when used between classes in different fragments. Such a use of class inheritance is similar to *mixin-style inheritance* [3, 12]: superclass implementations are not known until linking. Mixins have well-known type-checking challenges related to ambiguous methods: a subclass may “introduce” a method that conflicts with a method already provided by a superclass at link-time.

Existing OO languages that support modular linking do so by restricting mixin-style inheritance, disallowing cyclic dependencies between fragments, or by severely limiting what can be hidden between fragments [2, 10, 16]. In this paper, we show how separate compilation and modular linking can be supported in an OO language without such restrictions on expressiveness. In previous work [18] we have demonstrated the value of removing these restrictions. For this paper, we use a small language called MiniJiazzi with a Java-like core language and a linking language based on program units [11]. MiniJiazzi forms the theoretical foundation for our previous work on Jiazzi [18], which enhances Java’s linking model with support for separately compiled and externally linked fragments.

We have included 5 pages of appendices that contain type rules and proof sketches, which the referees should feel free to ignore. Ignoring those pages, the paper consists of less than 11 pages of text and figures and 1 page of bibliography.

```

fragment base {
  class BaseWidget extends Object { Object paint() { return ... } }
  class BaseButton extends Widget {...}
}
fragment win32 {
  class Widget extends BaseWidget {
    Object winpaint(int x) { return ... }
    Object paint() { return ... }
  }
  class Button extends BaseButton { Object winpaint(int x) { return ... } }
}
fragment motif {
  class Widget extends BaseWidget {
    Object motifpaint(int x) { return ... }
    Object paint() { return ... }
  }
  class Button extends BaseButton { Object motifpaint(int x) { return ... } }
}

```

Figure 1: A set of fragments that implement either Win32 or Motif UI widget classes, depending on whether the fragment base is linked against the win32 or motif fragments.

Issues related to the combination of modular linking and object-oriented languages have been previously explored in the areas of managing virtual method namespaces [19, 22], merging module systems and OO languages [2, 10, 16], and reasoning about mixin-style inheritance [3, 12]. The work presented in this paper is the first to show how to add modular linking to an OO language and support both cyclic linking and expressive mixin-style inheritance. Our type-checking architecture is also novel. Modular type checking in MiniJiazi ensures that compile-time type checking and link-time type checking environments are **congruent** with an intermediate type checking environment that does not change between compile-time and link-time. Defining such congruence is key to the correctness and expressiveness of the linking language.

Section 2 briefly motivates modular linking for an OO language and describes the challenges that must be overcome when adding modular linking to an OO language. Section 3 introduces the MiniJiazi language and describes our type checking methodology for adding modular linking to an OO language. Section 4 shows how we deal with cyclic dependencies between fragments in the presence of mixin-style inheritance. Section 5 discusses related work and Section 6 summarizes our conclusions.

2 Motivation

We consider a *fragment* to be a collection of classes compiled together into a single binary. Linking combines multiple fragments into a program that is a collection of all of the classes in the fragments. In OO languages like Java, there is no explicit support for fragments; instead, they are defined indirectly by a developer using constructs such as JAR files, codebases, and packages. In this section, we use examples from Java and explicitly define logical fragment boundaries. For the rest of this paper, we use Java as our prototypical OO language.

Fragments in Java can be linked with fragments that they were not compiled with, subject to certain binary compatibility [7, 13] rules. Binary compatibility supports fragment evolution: it allows an older application to be linked with a newer version of a class library, and it also supports component software development. Consider the three fragments in Figure 1 that implement parts of user interface (UI) widget classes. The base fragment provides a base implementation of UI widget classes while the win32 and motif fragments add platform-specific implementations to the UI widget classes. Linking the base fragment with either the win32 or motif fragments (but not both) yields a configuration of the UI widget classes that is appropriate for a specific platform.

The example in Figure 1 demonstrates the utility of mixin-style inheritance across fragment boundaries and cyclic linking of fragments. A special inheritance pattern is used in the fragment base to ensure that the button class is a subclass of all the widget classes. Rather than having the class BaseButton inherit directly from class BaseWidget, it in-

```

fragment person_v1 {
  class Person extends Object {}
}
fragment person_v2 {
  class Person extends Icon {}
}
fragment cowboyicon {
  class Cowboy extends Person {
    Object duel() { return this.draw(); }
    Object draw() { return ... /* cowboys draw guns */ }
  }
  class Icon extends Object {
    Object paint() { return this.draw(); }
    Object draw() { return ... /* draw an icon */ }
  }
}

```

Figure 2: The program fragment `cowboyicon` is originally compiled against `person_v1` but is actually linked against program fragment `person_v2`.

herits directly from the class `Widget` provided by one of the platform-specific fragments, which in turn inherits from the class `BaseWidget`—an exported class inherits from an imported class! This technique allows enhancements to the widget class to also enhance the button class by using a cyclic dependency between the `base` fragment and either of the platform-specific fragments. The fragments in Figure 1 allow methods and fields to be modularly added to existing classes. This usage of fragments is a static alternative to design patterns such as abstract factories [14], that support dynamic configuration of feature sets. Our previous work [18] discusses the benefits of and mechanics behind this usage of fragments in Jiazzi at greater length.

In normal Java, the linking in Figure 1 can be implemented by creating a custom classloader to implement the required linking logic. However, such linking is fragile, because Java does not support modular linking; it depends on global type checking after link-time to ensure global program type correctness [7, 17]. Global type checking complicates dynamic loading and makes it difficult to debug linking errors. Java’s global checking also cannot detect method name ambiguities, because its class signatures do not contain enough information.

Cyclic linking of fragments and mixin-style inheritance of classes between fragments complicates modular type checking: in particular detection of method name ambiguities. When a method in a superclass has the same name as a “new” method in a subclass, a linker could either resolve the ambiguity automatically or reject the mixin instantiation as a method collision. Automatically resolving the ambiguity can cause separate development of components to go awry, as shown by the three fragments in Figure 2. The fragment `cowboyicon` is originally compiled against the fragment `person_v1`, but it is actually linked into a program with fragment `person_v2`. When the programmer developed the fragment `cowboyicon`, the class `Cowboy` was not a subclass of the class `Icon`, so the method `draw` in class `Cowboy` is not intended to override the method `draw` in class `Icon`. However, an inheritance relationship between the classes `Person` and `Icon` is created in `person_v2` that did not exist in `person_v1`, and as a result `Cowboy` becomes a subclass of `Icon`, creating cowboy icons that draw their guns when painted! Although the program created by linking the fragments `cowboyicon` and `person_v2` is correctly typed, the ambiguity in Figure 2 cannot be detected in Java, because Java is not designed to detect (and reject) method collisions of this sort.

3 MiniJiazzi

MiniJiazzi consists of a small Java-like core language, which is used to implement the classes in fragments, and a separate linking language, which is used to link fragments together into a program. The MiniJiazzi core language is independent from the linking language with its own separate properties, which allows us to modularize our reasoning and soundness proof of the entire MiniJiazzi language. The separation of the core language and linking language is also reflected in Jiazzi [18]; Jiazzi is implemented as an enhancement to Java and does not require changes to the Java core language or its core type checking rules. We forego an detailed discussion of the MiniJiazzi core language

so that the abstract relationship between the linking language and the core language can be emphasized. A detailed description of the MiniJiazzi core language, whose syntax, type checking, and evaluation are similar to other small Java-like languages [12, 15], is given in Appendix A.

Notation: Uppercase letters identify constructs: e.g., A, B, and C identify classes and U, V, and W identify fragments; lowercase letters (possibly subscripted) designate constructs: e.g., c_a designates a class expression and s_b designates a class signature. Rules are in small caps, e.g., RULE-OK. Directed overbars ($\overrightarrow{\quad}$) are vectors that designate a sequence, e.g., \overrightarrow{A} is a sequence of class identifiers and \overrightarrow{u} is a sequence of unit expressions. Multiple sequences can be joined using the union operator (\cup). A rule within a vector applies to all elements specified by the vector. Terms adjacent to the rule but outside of the vector are copied as the vector is expanded. For example, if $\overrightarrow{A} = B C$ then $D \text{ RULE-OK } \overrightarrow{A}$ expands to $D \text{ RULE-OK } B$ and $D \text{ RULE-OK } C$.

3.1 Core Language

The MiniJiazzi core language is used to construct class expressions (designated by “c”), which are similar to Java class definitions: they express subclassing relationships with other classes and implement methods (which possibly override methods defined in super classes). Compilation and type checking of a class expression requires a typing environment, which describes the structure of the class expression’s dependencies. While these dependencies are other class expressions, a typing environment is not formed directly from class expressions. Instead a typing environment can be formed from *class signatures* (designated by “s”), which describe the subclassing relationships and available methods of corresponding class expressions without exposing implementation details. Finally, class signatures must declare what “new” methods they introduce, so that mixin constructions can be checked appropriately.

Type checking of a class expression ensures that a class expression is well-formed with respect to a well-formed typing environment. A typing environment is well-formed if the following criteria are ensured: first, a typing environment must be closed: no classes can be referenced by the typing environment unless they are described in the typing environment; second, there are no inheritance cycles expressed within the typing environment: all classes are direct or indirect subclasses of `Object`; third, there are no method collisions within the typing environment: no method is introduced by a subclass if that method is already introduced by a superclass.

A class expression is well-formed with respect to a typing environment if its methods are implemented correctly according to the typing environment: e.g., instantiated classes and invoked methods must be described by the typing environment. The rule ENVIRONMENT-OK ensures that a typing environment is well-formed, and the rule EXPRESSION-OK ensures that a class expression is well-formed in the presence of a typing environment. Without considering the linking language, a program in the MiniJiazzi core language is a sequence of class expressions (\overrightarrow{c}) that can be type checked together to ensure program type correctness using the rule GLOBAL-OK, which is defined in the following judgement:

$$\frac{\begin{array}{c} \vdash_{tc} \text{ENVIRONMENT-OK } \overrightarrow{c} \quad \overrightarrow{c} \quad \vdash_{tc} \text{EXPRESSION-OK } \overrightarrow{c} \\ \hline \vdash_{tc} \text{GLOBAL-OK } \overrightarrow{c} \end{array}}$$

The typing environment is formed by combining the class signatures from all class expressions of the program; the signature extraction operator returns the signature of a class (where $\| c \| = s$). GLOBAL-OK is not a modular rule, and is only presented here as a point of comparison. We assume that type checking described by the rule GLOBAL-OK and standard evaluation rules are sound, where a soundness proof is similar to other Java-like small languages [12, 15].

3.2 Linking Language

In order for MiniJiazzi to support modular linking, it must support fragments that are subject to deployment and composition into a program. In MiniJiazzi, a fragment is realized as a *unit expression* (designated by “u”). A unit expression has a *signature* (designated by “v”), which describes the unit’s interface with other units, and an *implementation*, which is a sequence of class expression. A unit signature consists of the signatures of imported classes that are

provided by other units and the signatures of exported classes that are provided to other units. Here is the syntax of a unit, and the definition of the signature extraction operator ($\| u \| = v$).

$$u ::= v_{sig} \{ \overrightarrow{c_{impl}} \} \quad v ::= U \{ \overrightarrow{s_{import}} \} \overrightarrow{s_{export}} \quad \| u = v_{sig} \{ \overrightarrow{c_{impl}} \} \| = v_{sig}$$

We use U , V , and W as unit identifiers. If type checking using the MiniJiazzi linking language is to be modular, type checking must be separated into two phases. *Compile-time* type checking ensures that a unit is internally consistent; it only examines one unit at a time. *Link-time* type checking ensures that units are externally consistent with each other: it ensures that all classes imported into linked units are properly satisfied by classes exported from linked units. Imported and exported classes are expressed in unit signatures, so unit implementations do not need to be examined during link-time type checking. A program in MiniJiazzi is a sequence of unit expressions (\vec{u}) that is type-checked to ensure program type correctness using the rule MODULAR-OK:

$$\frac{\vdash_{tc} \overrightarrow{\text{COMPILE-OK}} \vec{u} \quad \vdash_{tc} \overrightarrow{\text{LINK-OK}} \| \vec{u} \|}{\vdash_{tc} \overrightarrow{\text{MODULAR-OK}} \vec{u}}$$

Where the rule COMPILER-OK implements compile-time type checking and the rule LINK-OK implements link-time type checking. Given a definition of *link reduction*, which is the process of transforming classes from their unlinked form within a unit to their linked form within a program, we can compare rules MODULAR-OK and GLOBAL-OK with Lemma 1:

Lemma 1 (MODULAR-IMPLIES-GLOBAL) *If a program of units is type correct using modular type-checking rules, then the link-reduced units, which together yield a program of classes, are type-correct using global type checking rules.*

$$\vdash_{tc} \overrightarrow{\text{MODULAR-OK}} \vec{u} \quad \| \vec{u} \| \vdash_l \vec{u} \rightarrow \vec{c} \quad \Rightarrow \quad \vdash_{tc} \overrightarrow{\text{GLOBAL-OK}} \vec{c}$$

Link reduction is performed using the \vdash_l judgement from a unit to its link-reduced classes. Lemma 1 can be proven given the appropriate definitions of the type-checking rules COMPILER-OK and LINK-OK, described in Section 3.3, and link reduction, described in Section 3.4.

3.3 Congruence

Global type checking of a monolithic program with the rule GLOBAL-OK uses only one typing environment; modular type checking of a program composed out of units consists of three kinds of typing environments. A *compilation typing environment* is used to perform compile-time type checking of a unit's implementation. A *linkage typing environment* is used to perform link-time type checking between units as they are linked into a complete program. Finally, an *intermediate typing environment* acts as a bridge between the compilation and linkage typing environments, so modular type checking can ensure global program type correctness.

Separate compilation is enabled in Jiazzi by having unit signatures contain the signatures of imported classes. To form the compilation typing environment of a unit, the class signatures extracted from the class expressions of the unit's implementation are combined with imported class signatures as $\| \overrightarrow{c_{impl}} \| \cup \overrightarrow{s_{import}}$. Since units are compiled separately, linking a program in MiniJiazzi must ensure that a unit's imports are adequately satisfied by classes exported from other unit's of the program. Unit signatures also embed the signatures of exported classes, which avoids examining unit implementations during linking. A program need only be relinked when the signature of a linked unit changes. The linkage typing environment is formed from the exported class signatures of every unit as $\overrightarrow{s_{export}}$.

Compile-time type checking of a unit is implemented in the rule COMPILER-OK, which ensures that each class expression contained within the unit is well-formed (using the core language type checking rule EXPRESSION-OK with respect to the unit's compilation typing environment). Link-time type checking of a unit is implemented in the rule LINK-OK. The intermediate typing environment is formed from the signatures of a unit's imported and exported

classes as $\overrightarrow{s_{import}} \cup \overrightarrow{s_{import}}$. All typing environments must be well-formed, which is ensured using the core language type checking rule ENVIRONMENT-OK described earlier.

The key to making type checking modular in MiniJiazzi is ensuring that the compilation typing environments of each linked unit and linkage typing environment are compatible. This can be achieved by ensuring that for each unit, the unit's compilation typing environment and the linkage typing environment are **congruent** with the unit's intermediate typing environment. The CONGRUENT rule enforces this relationship, and is used to define the rules COMPILE-OK and LINK-OK:

$$\frac{\begin{array}{c} \vdash_{tc} \text{ENVIRONMENT-OK } \overrightarrow{s_i} \cup \overrightarrow{c} \parallel \overrightarrow{c} \parallel \quad \overrightarrow{s_i} \cup \overrightarrow{c} \parallel \vdash_{tc} \overrightarrow{\text{EXPRESSION-OK } c} \\ \vdash_{tc} \text{ENVIRONMENT-OK } \overrightarrow{s_i} \cup \overrightarrow{s_e} \quad \vdash_{tc} \overrightarrow{s_i} \cup \overrightarrow{c} \parallel \text{CONGRUENT } \overrightarrow{s_i} \cup \overrightarrow{s_e} \end{array}}{\vdash_{tc} \text{COMPILE-OK } u = U \{ \overrightarrow{s_i} \overrightarrow{s_e} \} \{ \overrightarrow{c} \}}$$

$$\frac{\begin{array}{c} \vdash \text{UNIQUE } \overrightarrow{U} \\ \vdash_{tc} \text{ENVIRONMENT-OK } \overrightarrow{s_e} \\ \vdash_{tc} \overrightarrow{s_e} \text{CONGRUENT } \overrightarrow{s_i} \cup \overrightarrow{s_i} \end{array}}{\vdash_{tc} \text{LINK-OK } v = U \{ \overrightarrow{s_i} \overrightarrow{s_e} \}}$$

Congruence ensures that the typing environment to the left of the rule CONGRUENT, known as the *source typing environment*, is substitutable with the typing environment on the right of the rule CONGRUENT, known as the *sink typing environment*. This substitutability ensures that modular type checking can ensure that the resulting *program typing environment*, which is formed from the results of all link-reduced units, is also well-formed. We state this as Lemma 2:

Lemma 2 (PROGRAM-ENVIRONMENT-WELL-FORMED) *If compilation, intermediate, and linkage typing environments are well-formed, compilation typing environments are congruent with intermediate typing environments, and the linkage typing environment is congruent with all intermediate typing environments, then the linked reduced classes of a unit program together constitute a well-formed program typing environment.*

$$\frac{\begin{array}{c} \overrightarrow{u} = v \{ \overrightarrow{c_x} \} \quad \overrightarrow{v} = U \{ \overrightarrow{s_i} \overrightarrow{s_e} \} \quad \vdash_{tc} \text{ENVIRONMENT-OK } \overrightarrow{s_i} \cup \overrightarrow{c_x} \parallel \quad \vdash_{tc} \overrightarrow{s_i} \cup \overrightarrow{c_x} \parallel \text{CONGRUENT } \overrightarrow{s_i} \cup \overrightarrow{s_e} \\ \vdash_{tc} \text{ENVIRONMENT-OK } \overrightarrow{s_i} \cup \overrightarrow{s_e} \quad \vdash_{tc} \overrightarrow{s_e} \text{CONGRUENT } \overrightarrow{s_i} \cup \overrightarrow{s_e} \quad \vdash_{tc} \text{ENVIRONMENT-OK } \overrightarrow{s_e} \end{array}}{\vdash \text{UNIQUE } \overrightarrow{U} \quad \overrightarrow{v} \vdash_l \overrightarrow{u} \rightarrow \overrightarrow{c_y} \quad \Rightarrow \quad \vdash_{tc} \text{ENVIRONMENT-OK } \overrightarrow{c_y} \parallel}$$

Proving Lemma 2 is a large part of the proof for Lemma 1. The rule CONGRUENT must at least ensure that the source typing environment provides every class, method, and subclassing relationship that exists in the sink typing environment. However, this weakest definition of congruence would not be sufficient to detect inheritance cycles and method collisions in the presence of cyclic linking. The strongest definition of CONGRUENT would ensure that the source typing environment is equivalent to the sink typing environment—that the source typing environment does not provide any class, method, or subclassing relationship that does not exist in the sink typing environment. While this strongest definition of congruence would be safe, it is too restrictive: every class, method, and subclassing relationship of every unit would have to be visible. Such a requirement would defeat the purpose of separate compilation and would prevent a unit from being reused in multiple programs. In Section 4 we define congruence in a way that is safe and also allows for sufficient information hiding.

3.4 Link Reduction

When units are combined together into a linked program, linking must ensure that name clashes between unit implementations do not cause ambiguities at runtime. Name clashes in OO languages result from classes from different units with the same name, or distinct methods with the same names that are defined in compatible classes that originate from different units. While modular type checking can detect name clashes between units if classes and methods are exposed in unit signatures, all but the strongest congruence conditions permit some hiding between units. Finally,

classes and methods hidden within a unit's implementation should not clash with classes and methods from other units. Link reduction ensures that references to hidden classes and methods are disambiguated during linkage.

MiniJiazzi achieves disambiguation with *linking offsets*. Every class and method reference within a unit implementation is qualified with a linking offset that is utilized during evaluation to disambiguate references. Method expressions also have linking offsets to ensure that they implement or override the appropriate method from a superclass. Linking offsets are not provided by the programmer; they are assigned by the linker, in much the same way as branch offsets are rewritten when a dynamically-linked library (DLL) is loaded into memory. Before link reduction occurs, all linking offsets of a unit's implementation are empty (type checking does not use linking offsets). Link reduction then binds linking offsets according to the unit that the class or method originates from. The following are judgments that determine how linking offsets are bound for class references:

$$\frac{A[o] \in \overrightarrow{|s_{iu}|} \quad A[o] \in \overrightarrow{|s_{ew}|} \quad W \overrightarrow{s_{ew}} \in \overrightarrow{V \overrightarrow{s_e}}}{v = V \{ \overrightarrow{s_i} \overrightarrow{s_e} \}, u = U \{ \overrightarrow{s_{iu}} \overrightarrow{s_{eu}} \} \{ \overrightarrow{c} \} \vdash_l A \mapsto W} \quad \frac{A[o] \notin \overrightarrow{|s_{iu}|}}{v = V \{ \overrightarrow{s_i} \overrightarrow{s_e} \}, u = U \{ \overrightarrow{s_{iu}} \overrightarrow{s_{eu}} \} \{ \overrightarrow{c} \} \vdash_l A \mapsto U}$$

Classes are always identified by the terms A, B, C, and D and an identifier extraction operator returns the name of a class from a class signature ($|s| = A[W]$). The symbol $[o]$ indicates an empty linking offset and linking offsets are always in brackets ($[U]$). The map operator (\mapsto) is used to find the appropriate linking offset for a class reference in a unit u . For a reference to an imported class, the linking environment is used to determine which unit an imported class is provided by, using that unit as the linking offset. The linking offsets of references to classes that are not imported into the referring unit are bound to the referring unit as these references are to classes contained within the unit. Method reference linking offsets are bound using a similar set of judgments:

$$\frac{\overrightarrow{s_{iu}} \cup \overrightarrow{|c|} \vdash_m \langle A[o].M[o] \rangle = B[o] \quad B[o] \in \overrightarrow{|s_{iu}|} \quad \overrightarrow{s_e} \vdash_m \langle B[o].M[o] \rangle = C[o] \quad \overrightarrow{s}, u \vdash_l C \mapsto W}{v = V \{ \overrightarrow{s_i} \overrightarrow{s_e} \}, u = U \{ \overrightarrow{s_{iu}} \overrightarrow{s_{eu}} \} \{ \overrightarrow{c} \} \vdash_l A.M \mapsto W} \quad \frac{\overrightarrow{s_{iu}} \cup \overrightarrow{|c|} \vdash_m \langle A[o].M[o] \rangle = B[o] \quad B[o] \notin \overrightarrow{|s_{iu}|}}{v = V \{ \overrightarrow{s_i} \overrightarrow{s_e} \}, u = U \{ \overrightarrow{s_{iu}} \overrightarrow{s_{eu}} \} \{ \overrightarrow{c} \} \vdash_l A.M \mapsto U}$$

The method introduction operator $\overrightarrow{s} \vdash_m \langle A[U].M[V] \rangle = B[W]$ determines that the class $B[W]$, a super class of $A[U]$, introduces the method $M[V]$ as expressed in the typing environment \overrightarrow{s} . If, according to the compilation typing environment, the class that introduces a method is not imported, then the linking offset of the method is bound to the referencing unit. Otherwise, the linking offset of a method reference is bound to the unit that exports the superclass that introduces the method, which can be located using the linkage typing environment.

The MiniJiazzi units shown in Figure 3 demonstrate how link reduction resolves method ambiguity. The unit `icon` specifies what it expects from other units through its imports. Since the unit `icon` does not import the method `draw` from the `cowboy` unit, within the implementation of class `Icon` link reduction binds the linking offset of the call to method `draw` to the unit `icon`. Since the unit `cowboy` exports the method `draw`, whereas the unit `icon` does not, linking reduction binds the linking offset of the call to `draw` within the class `Main` to the unit `cowboy`.

In Java, method scope is established by packages and access flags. If the above example was written in normal Java using packages and access flags (the method `draw` in a package `cowboy` would be public, and the method `draw` in a package `icon` would be private), ambiguity between the `draw` methods could not be avoided and the Java source compiler would reject such a construction. Using unit signatures to establish the namespace of a unit makes MiniJiazzi more flexible than Java, because a unit can restrict its imported namespace as well as its exported namespace.

Rather than use linking offsets, other approaches [19, 20, 22] disambiguate between methods using dictionaries that are based on scope. Linking offsets simulate dictionaries by using automatic renaming of class and method references. MiniJiazzi avoids the use of dictionaries during evaluation, which allows us to reason more easily about the differences between the linking language and the pure core language: we do not require a separate set of evaluation rules for programs that use the linking language.


```

unit cowboy {
  import
  export class Cowboy extends Object { Object draw(), Object dual() }
} {
  class Cowboy[cowboy] extends Object {
    introduces Object draw()
  } {
    Object draw[cowboy]() { .../* draw guns in a duel */ },
    Object dual[cowboy]() { this.draw[cowboy]() }
  }
}
unit icon {
  import class Cowboy extends Object { }
  export class Icon extends Cowboy { Object paint() }
} {
  class Icon[icon] extends Cowboy[cowboy] {
    introduces Object paint(), Object draw()
  } {
    Object paint[icon]() { this.draw[icon]() },
    Object draw[icon]() { .../* draw an icon */ }
  }
}
unit main {
  import class Icon extends Object { Object paint(), Object draw() }
  export class Main extends Object { Object main() }
} {
  class Main[main] extends Object {
    introduces Object main()
  } {
    Object main[main]() { this.draw[cowboy]() }
  }
}

```

Figure 3: Three linked units in MiniJiazzi that demonstrate how linking offsets eliminate ambiguity in the *draw* method for classes *Cowboy* and *Icon*.

Link reduction only binds class and method reference linking offsets; it does not otherwise change the structure of class expressions within a unit. The weakest congruence conditions ensure that referenced classes and methods exist within the program typing environment and that method implementations override methods correctly. As a result, showing that the classes that result from link-reduced unit expressions are well-formed with respect to the program typing environment is only a matter of structural induction. The only nontrivial aspect of proving Lemma 1 is proving Lemma 2. We concentrate on the proof, which depends on the proper definition of the rule CONGRUENT, in Section 4.

4 Congruence

Congruence between typing environments is complicated by the fact that fragments can be linked in cycles. With cyclic linking, the weakest form of congruence (where a congruent typing environment is only required to have every class, method, and subclassing relationship expressed by the target environment) is insufficient to ensure global program type correctness. The strongest form of congruence is not desirable either, as it prevents any hiding between fragments. Finding a definition of congruence “in the middle” is complicated by the fact that class inheritance can span fragment boundaries: allowing a congruent typing environment to provide too many methods or subclassing relationships could lead to inheritance cycles or method collisions. Supporting cyclic dependencies between fragments is also problematic in many other modular languages (e.g., ML [6]).

For OO languages, we have discovered an effective compromise between the weakest and strongest congruence conditions. In addition to providing every subclassing relationship of the target typing environment, a source typing environment must express only subclassing relationships that do not conflict with subclassing relationships of the sink typing environment. If two classes in the sink typing environment do not have a subclassing relationship, then those

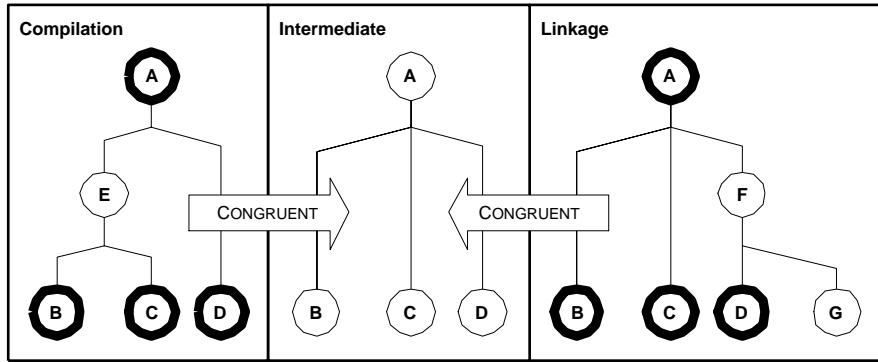


Figure 4: Three class-inheritance hierarchies expressed by compilation, intermediate, and linkage typing environments; classes in the compilation and linkage typing environments that are also in the intermediate typing environment are in bold; congruence conditions are indicated as arrows.

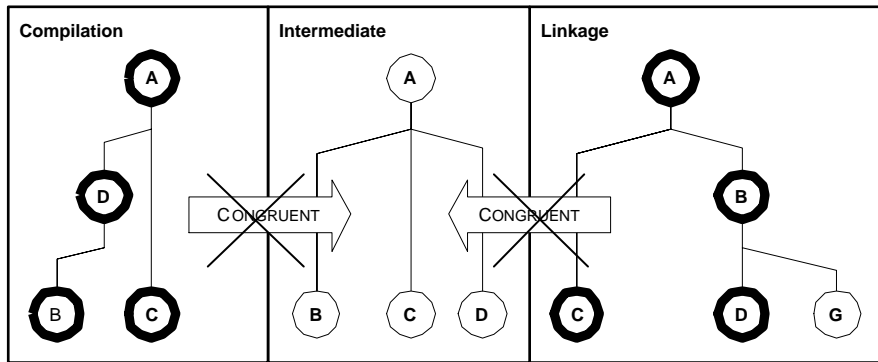


Figure 5: Examples of compilation and linkage typing environments that are not congruent with the intermediate typing environment.

two classes in the source typing environment must also not have a subclassing relationship. This congruence condition allows classes and methods to be hidden between fragments and still allows modular type checking to detect method collisions and inheritance cycles.

In Figure 4, both the compilation and linkage typing environments are congruent with the intermediate typing environment. The intermediate typing environment contains the signatures for the classes A, C, D, and E, where D, C, and E are subclasses of A. The compilation typing environment contains the signature for class B in addition to classes contained in the intermediate typing environment. The compilation typing environment is congruent with the intermediate typing environment because classes C and D are also subclasses (albeit indirect) of A in the compilation typing environment, and no new subclassing relationship between A, C, D, and E is expressed in the compilation typing environment that is not already expressed in the intermediate typing environment. A similar argument shows that the linkage typing environment is congruent with the intermediate typing environment.

In Figure 5, neither the compilation nor the linkage typing environments are congruent with the intermediate typing environment. While each typing environment expresses every subclassing relationship expressed by the intermediate typing environment, each establishes a subclassing relationship incompatible with the intermediate typing environment. Both the compilation and linkage typing environments express a subclassing relationship between classes B and D, while the intermediate typing environment expresses no subclassing relationship between these two classes. As a result, if program linking were permitted, an inheritance cycle would be created between classes B and D.

In MiniJazzi, the rule CONGRUENT takes the following form:

$$\begin{array}{c}
\frac{\overrightarrow{\mathfrak{s}}_{source} \vdash_t A[o] \triangleleft \mathfrak{t}_s}{\overrightarrow{\mathfrak{s}}_{source}, \overrightarrow{\mathfrak{s}}_{sink} \vdash_{tc} A \text{ SUPER-OK } \mathfrak{t}_s} \\
\\
\frac{\overrightarrow{\mathfrak{s}}_{source} \vdash_t A[o] \triangleleft B[o] \quad B[o] \notin |\overrightarrow{\mathfrak{s}}_{sink}|}{\overrightarrow{\mathfrak{s}}_{source}, \overrightarrow{\mathfrak{s}}_{sink} \vdash_{tc} B \text{ SUPER-OK } \mathfrak{t}_s} \\
\overrightarrow{\mathfrak{s}}_{source}, \overrightarrow{\mathfrak{s}}_{sink} \vdash_{tc} A \text{ SUPER-OK } \mathfrak{t}_s \\
\\
\frac{\overrightarrow{\mathfrak{s}}_{source} \vdash_m \langle\langle A[o] \rangle\rangle = \overrightarrow{\mathfrak{n}}_{source} \quad \overrightarrow{\mathfrak{s}}_{sink} \vdash_m \langle\langle A[o] \rangle\rangle = \overrightarrow{\mathfrak{n}}_{sink} \quad |\overrightarrow{\mathfrak{s}}_{source}| \supseteq |\overrightarrow{\mathfrak{s}}_{sink}| \quad \overrightarrow{\mathfrak{n}}_{source} \supseteq \overrightarrow{\mathfrak{n}}_{sink} \quad \overrightarrow{\mathfrak{s}}_{source}, \overrightarrow{\mathfrak{s}}_{sink} \vdash_{tc} A \text{ SUPER-OK } \mathfrak{t}_s}{\vdash_{tc} \overrightarrow{\mathfrak{s}}_{source} \text{ CONGRUENT } \overrightarrow{\mathfrak{s}}_{sink} = A[o] \triangleleft \mathfrak{t}_s \{ \overrightarrow{\mathfrak{n}} \}}
\end{array}$$

The method operator $\overrightarrow{\mathfrak{s}} \vdash_m \langle\langle A[U] \rangle\rangle = \overrightarrow{\mathfrak{n}}$ extracts the sequence of all methods in the class $A[U]$ and its superclasses, given the typing environment $\overrightarrow{\mathfrak{s}}$. The rule CONGRUENT can enforce the congruent subclassing relationship between a source typing environment and a sink typing environment using a recursively defined rule SUPER-OK. Intuitively, SUPER-OK ensures that direct subclassing relationships in the sink environment are satisfied with (possibly indirect) subclassing relationships in the source environment without first proceeding through any classes also present in the sink environment (otherwise, SUPER-OK is stuck). The rule CONGRUENT permits method hiding and some hiding of subclass relationships, and still enables the detection of method collisions and inheritance cycles.

4.1 Proof Sketch

The SUPER-OK rule ensures that all subclassing relationships visible in a sink typing environment are visible in the source typing environment, and that the source typing environment does not introduce any subclassing relationships that conflict with the sink typing environment. We can express this as Lemma 3:

Lemma 3 (SUPER-OK-EQUIVALENCE) *With respect to classes visible in a sink typing environment, the source typing environment and sink typing environment agree on the subclassing relationships between these classes.*

$$\begin{array}{l}
\overrightarrow{\mathfrak{s}}_{sink} = \overrightarrow{\mathfrak{C}[o]} \triangleleft \mathfrak{t}_s \dots \quad |\overrightarrow{\mathfrak{s}}_{source}| \supseteq |\overrightarrow{\mathfrak{s}}_{sink}| \quad \overrightarrow{\mathfrak{s}}_{source}, \overrightarrow{\mathfrak{s}}_{sink} \vdash_t \overrightarrow{\mathfrak{C}} \text{ SUPER-OK } \mathfrak{t}_s \\
\Rightarrow \quad \forall A[o], B[o] \in \overrightarrow{\mathfrak{s}}_{sink} \quad \overrightarrow{\mathfrak{s}}_{sink} \vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{\mathfrak{s}}_{source} \vdash_t A[o] < B[o]
\end{array}$$

The proof of Lemma 2 depends primarily on Lemma 4:

Lemma 4 (LOCAL-EQUIVALENCE) *With respect to classes visible in a compilation typing environment, the compilation typing environment and linkage typing environment agree on the subclassing relationships between these classes.*

$$\begin{array}{l}
\vdash_{tc} \text{MODULAR-OK } \mathfrak{u} = \mathfrak{V} \{ \overrightarrow{\mathfrak{s}}_i \overrightarrow{\mathfrak{s}}_e \} \{ \overrightarrow{\mathfrak{c}}_x \} \quad \|\mathfrak{u}\| \vdash_l \mathfrak{u} \rightarrow \overrightarrow{\mathfrak{c}}_y \quad \Rightarrow \\
\forall \mathfrak{u}_u \in \overrightarrow{\mathfrak{u}} \quad \mathfrak{u}_u = \mathfrak{U} \{ \overrightarrow{\mathfrak{s}}_{iu} \overrightarrow{\mathfrak{s}}_{eu} \} \{ \overrightarrow{\mathfrak{c}}_{xu} \}, \quad \forall A[o], B[o] \in \overrightarrow{\mathfrak{s}}_{iu} \cup \|\mathfrak{c}_{xu}\| \\
\|\mathfrak{u}\|, \mathfrak{u}_u \vdash_l A \mapsto \mathfrak{V}, B \mapsto \mathfrak{W} \quad \overrightarrow{\mathfrak{s}}_{iu} \cup \|\mathfrak{c}_{xu}\| \vdash_t A[o] < B[o] \leftrightarrow \|\mathfrak{c}_y\| \vdash_t A[\mathfrak{V}] < B[\mathfrak{W}]
\end{array}$$

A proof for Lemma 4 that primarily depends on Lemma 3 is shown in Appendix B. We can prove Lemma 2 by showing that inheritance cycles and method collisions are impossible. The absence of inheritance cycles falls directly out of Lemma 4, as the compilation typing environment of each unit does not have cycles, so the resulting link-reduced classes cannot express any cycles either. Method collisions can only occur between methods introduced in the same unit, because link reduction specifies different linking offsets to methods that originate from different units. If a method collision does occur, it must be created within one unit, where it must have been detectable according to the subclassing relationships ensured by Lemma 4.

5 Related Work

Separate compilation and modular linking are explored by Cardelli [4] in his work with linksets. Program units [8, 11] extend work in this area by considering higher-order core languages, circular dependencies, and dynamically typed classes. MiniJiazzi shows how program units can be added to statically-typed object-oriented languages.

Drossopoulou et al. [7] point out problems in Java’s linking model, especially the fact that is not very safe, and propose a safer model for linking in Java. Ancona et al. [1] use a notion of a compilation schema to explore separate source code compilation and runtime linking in Java as it is currently. Our work proposes a new compilation and linking model for OO languages that addresses the deficiencies in Java’s existing compilation and linking models.

Ancona and Zucca [2] explore adding a true module to the Java language with a language called JavaMod. Fisher and Reppy [9, 10] explore modular linking for a class-based core language using a ML-like module system called Moby. Unlike MiniJiazzi, neither JavaMod nor Moby support true mixin-style inheritance. In JavaMod, methods hidden in a superclass are not visible in a subclass, while in Moby, methods provided by a superclass can only be invoked by explicitly specifying the superclass. Moby does not support modules with cyclic dependencies. As far as module systems go, MiniJiazzi is the first formalization of a statically-typed module system that supports both cyclic dependencies and full mixin-style inheritance for an OO language.

Enforcing the privacy of methods in OO languages has been explored extensively in the literature. Riecke and Stone [19] formally explore method privacy in structurally typed OO languages by using method dictionaries, while this work is extended by Stone [20] and Vouillon [22] in the context of class and mixin-based languages. MiniJiazzi differs by using linking offsets rather than dictionaries to enforce method scopes and disambiguate between method namespaces.

6 Conclusion

We have shown how modular linking in an OO language can be supported while allowing several expressive features: cyclic dependencies between fragments, inheritance across fragment boundaries, and information hiding between fragments. MiniJiazzi forms the theoretical basis for our implementation of Jiazzi, a component system for the Java language. Jiazzi is more sophisticated than MiniJiazzi, but their type-checking rules are similar. An implementation of Jiazzi is available for download:

<http://www.cs.utah.edu/plt/jiazzi>

Not all features of OO languages facilitate modular linking. We did not have room to describe how abstract methods makes modular linking difficult: they limit information hiding between fragments. If a class has an abstract method, then that abstract method can never be hidden when the class is visible in another fragment. While abstract methods can be used in Jiazzi units, they are not as modular as normal virtual methods. Future work will further explore interactions of modular linking with such OO language features.

References

- [1] D. Ancona, G. Lagorio, and E. Zucca. A formal framework for Java separate compilation. In *In Proc. of ECOOP*, pages 609–636, June 2002.
- [2] D. Ancona and E. Zucca. True modules for Java classes. In *In Proc. of ECOOP*, pages 354–380, June 2001.
- [3] G. Bracha and W. Cook. Mixin-based inheritance. In *Proc. of OOPSLA*, pages 303–311, Oct. 1990.
- [4] L. Cardelli. Program fragments, linking and modularization. In *Proc. of POPL*, pages 266–277, Jan. 1997.
- [5] C. Clifton, G. Leavens, C. Chambers, and T. Millstein. MultiJava: Modular open classes and symmetric multiple dispatch for Java. In *Proc. of OOPSLA*, pages 130–146, Oct. 2000.
- [6] K. Crary, R. Harper, and S. Puri. What is a recursive module? In *Proc. of PLDI*, pages 50–63, May 1999.
- [7] S. Drossopoulou, D. Wragg, and S. Eisenbach. What is Java binary compatibility? In *Proc. of OOPSLA*, pages 341–361, Oct. 1998.

- [8] R. Findler and M. Flatt. Modular object-oriented programming with units and mixins. In *Proc. of ICFP*, pages 98–104, Sept. 1998.
- [9] K. Fisher and J. Reppy. The design of a class mechanism for Moby. In *Proc. of PLDI*, pages 37–49, May 1999.
- [10] K. Fisher and J. Reppy. Extending Moby with inheritance-based subtyping. In *Proc. of ECOOP*, pages 83–107, June 2000.
- [11] M. Flatt and M. Felleisen. Units: Cool modules for HOT languages. In *Proc. of PLDI*, pages 236–248, May 1998.
- [12] M. Flatt, S. Krishnamurthi, and M. Felleisen. Classes and mixins. In *Proc. of POPL*, pages 171–183, Jan. 1999.
- [13] I. Forman, M. Conner, S. Danforth, and I. Raper. Release-to-release binary compatibility in SOM. In *Proc. of OOPSLA*, Oct. 1995.
- [14] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1994.
- [15] A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In *Proc. of OOPSLA*, pages 132–146, Oct. 1999.
- [16] X. Leroy, D. Doligez, J. Garrigue, D. R’emy, and J. Vouillon. The Objective CAML system, documentation and user’s manual, 2000. <http://caml.inria.fr/ocaml/htmlman/>.
- [17] S. Liang and G. Bracha. Dynamic class loading in the Java Virtual Machine. In *Proc. of OOPSLA*, Oct. 1998.
- [18] S. McDirmid, M. Flatt, and W. Hsieh. Jiazzzi: New-age components for old-fashioned Java. In *Proc. of OOPSLA*, pages 211–222, Oct. 2001.
- [19] J. G. Riecke and C. A. Stone. Privacy via subsumption. *Theory and Practice of Object Systems*, 1999.
- [20] C. Stone. Extensible objects without labels. In *Proc. of FOOL*, Jan. 2002.
- [21] C. Szyperski. Import is not inheritance: Why we need both: modules and classes. In *Proc. of ECOOP*, pages 19–32, 1992.
- [22] J. Vouillon. Combining subsumption and binary methods: An object calculus with views. In *Proc. of POPL*, pages 290–303, Jan. 2001.

A MiniJiazzi Core Language

MiniJiazzi core language syntax

class-identifier – A, B, C, D method-identifier – M, N fragment-identifier – U, V, W, o
 class-signature s ::= A[U] < t_s { n_{fresh} } method-signature n ::= t_{return} M[U](t_{arg} x)
 class-expression c ::= s_{sig} { m_{impl} } method-expression m ::= n_{sig} { e_{return} }
 type t ::= Object | A[U]
 method-subexpression e ::= new A[U] | (t) e | x | this | e.M[A[U], V](e_x)

The syntax of the MiniJiazzi core language is a small Java-like language subset. To emphasize inheritance and virtual methods, the MiniJiazzi core language class expressions only express subclassing relationships and contain methods. The symbol < is used to express a direct subclassing relationship. As mentioned in Section 3.4, link offsets are added to the syntax of the core language even though programmers do not use them directly.

MiniJiazzi core language extraction, subtyping, and method relationships

$\| c = s \{ \overrightarrow{m_{impl}} \} \| = s$ $\| m = n \{ e \} \| = n$ $| A[U] < t_s \{ \overrightarrow{n} \} | = A[U]$ $| t_r M[U](\overrightarrow{t_x} \overrightarrow{x}) | = M[U]$

$$\frac{A[U] < t_s \{ \dots \} \in \overrightarrow{s}}{\overrightarrow{s} \vdash_t A[U], A[U] < t_s, A[U] < t_s} \quad \frac{\overrightarrow{s} \vdash_t t_a < t_b}{\overrightarrow{s} \vdash_t t_a \leq t_b} \quad \frac{\overrightarrow{s} \vdash_t t_a \leq t_a}{\overrightarrow{s} \vdash_t \text{Object}} \quad \frac{\overrightarrow{s} \vdash_t t_a \leq t_b < t_c}{\overrightarrow{s} \vdash_t t_a < t_c}$$

$$\overrightarrow{s} \vdash_m \langle \langle \text{Object} \rangle \rangle = \emptyset$$

$$\frac{A[U] < t_s \{ \overrightarrow{n_x} \} \in \overrightarrow{s} \quad \overrightarrow{s} \vdash_m \langle \langle t_s \rangle \rangle = \overrightarrow{n_y}}{\overrightarrow{s} \vdash_m \overrightarrow{n_x} \cup \langle \langle A[U] \rangle \rangle = \overrightarrow{n_y}}$$

$$\frac{A[U] < t_s \{ \overrightarrow{n} \} \in \overrightarrow{s} \quad \overrightarrow{s} \vdash_m M[U] \in \overrightarrow{n}}{\overrightarrow{s} \vdash_m \langle A[U].M[U] \rangle = A[U]}$$

$$\frac{A[U] < t_s \{ \overrightarrow{n} \} \in \overrightarrow{s} \quad M[W] \notin \overrightarrow{n} \quad \overrightarrow{s} \vdash_t A[U] < B[V] \quad \overrightarrow{s} \vdash_m \langle B[V].M[W] \rangle = C[W]}{\overrightarrow{s} \vdash_m \langle A[U].M[W] \rangle = C[W]}$$

MiniJiazzi core language type checking

$$\frac{\overrightarrow{s} \vdash_{tc} A[U] \overrightarrow{\text{METHOD-OK}} m \quad \overrightarrow{n} \subseteq \overrightarrow{\| m \|} \quad \vdash \text{UNIQUE} \overrightarrow{\| m \|}}{\overrightarrow{s} \vdash_{tc} \text{EXPRESSION-OK} \quad A[U] < t_s \{ \overrightarrow{n} \} \{ \overrightarrow{m} \}}$$

$$\frac{\vdash \text{UNIQUE} A[U] \quad \overrightarrow{s} \vdash_t \overrightarrow{t_x} \quad \overrightarrow{s} \vdash_t \overrightarrow{t_r} \quad \vdash \text{UNIQUE} M[U] \quad \overrightarrow{s} \vdash_t t_s \leq \text{Object} \quad \overrightarrow{s} \vdash_m \overrightarrow{n} \cap \langle \langle t_s \rangle \rangle = \emptyset}{\vdash_{tc} \text{ENVIRONMENT-OK} \quad s = A[U] < t_s \{ n = t_r M[U](\overrightarrow{t_x}) \}}$$

MiniJiazzi method-level type checking

$$\frac{\vec{s} \vdash_t \vec{t}_x \quad \vec{s} \vdash_t \mathbf{t}_r \quad \vdash \mathbf{UNIQUE} \vec{x}}{\Gamma(\mathbf{x}) = \vec{t}_x \quad \Gamma(\mathbf{this}) = \mathbf{A}[\mathbf{U}] \quad \Gamma, \vec{s} \vdash_{tc} e \in \mathbf{t}_y \quad \Gamma, \vec{s} \vdash_{tc} \mathbf{x} \in \Gamma(\mathbf{x}) \quad \frac{\vec{s} \vdash_t \mathbf{A}[\mathbf{U}]}{\Gamma, \vec{s} \vdash_{tc} \mathbf{new} \mathbf{A}[\mathbf{U}] \in \mathbf{A}[\mathbf{U}]}}{\vec{s} \vdash_{tc} \mathbf{A}[\mathbf{U}] \mathbf{METHOD-OK} \mathbf{t}_r \mathbf{M}[\mathbf{W}](\vec{t}_x \vec{x}) \{ e \}}$$

$$\frac{\Gamma, \vec{s} \vdash_{tc} e \in \mathbf{t}_b \quad \vec{s} \vdash_t (\mathbf{t}_a \leq \mathbf{t}_b \text{ or } \mathbf{t}_b < \mathbf{t}_a)}{\Gamma, \vec{s} \vdash_{tc} (\mathbf{t}_a) e \in \mathbf{t}_a} \quad \frac{\Gamma, \vec{s} \vdash_{tc} e \in \mathbf{t}_b \quad \vec{s} \vdash_t \mathbf{t}_a \not\leq \mathbf{t}_b, \mathbf{t}_b \not< \mathbf{t}_a}{\Gamma, \vec{s} \vdash_{tc} (\mathbf{t}_a) e \in \mathbf{error}} \quad \frac{\Gamma, \vec{s} \vdash_{tc} e \in \mathbf{t}, \overrightarrow{\mathbf{e}_x \in \mathbf{t}_y} \quad \vec{s} \vdash_t \mathbf{t} \leq \mathbf{A}[\mathbf{V}], \overrightarrow{\mathbf{t}_y \leq \mathbf{t}_x} \quad \vec{s} \vdash_m \mathbf{t}_r \mathbf{M}[\mathbf{W}](\vec{t}_x) \in \langle \langle \mathbf{A}[\mathbf{V}] \rangle \rangle}{\Gamma, \vec{s} \vdash_{tc} e.\mathbf{M}[\mathbf{A}[\mathbf{V}], \mathbf{W}](\overrightarrow{\mathbf{e}_x}) \in \mathbf{t}_r}$$

Method invocations specify a class offset in addition to a linking offset. This class offset is used only as a convenience to the linker and is not used during evaluation. Like link offsets, class offsets are not specified directly by the programmer. Rather they can be provided by the type checker during compile-time type checking.

MiniJiazzi core language evaluation

$$\frac{\vec{c} \vdash_c e \rightarrow \mathbf{new} \mathbf{A}[\mathbf{V}] \quad \|\vec{c}\| \vdash_t \mathbf{A}[\mathbf{V}] \leq \mathbf{t}}{\vec{c} \vdash_c (\mathbf{t}) e \rightarrow \mathbf{new} \mathbf{A}[\mathbf{V}]} \quad \frac{\vec{c} \vdash_c e \rightarrow \mathbf{new} \mathbf{A}[\mathbf{V}] \quad \|\vec{c}\| \vdash_t \mathbf{A}[\mathbf{V}] \not\leq \mathbf{t}}{\vec{c} \vdash_c (\mathbf{t}) e \rightarrow \mathbf{downcast_error}}$$

$$\frac{\vec{c} \vdash_c e \rightarrow \mathbf{new} \mathbf{A}[\mathbf{W}], \overrightarrow{\mathbf{e}_x \rightarrow \mathbf{new} \mathbf{t}_x} \quad \vec{c} \vdash_c [\mathbf{A}[\mathbf{W}].\mathbf{M}[\mathbf{V}]] = (\vec{x}) \{ e \}}{\vec{c} \vdash_c e.\mathbf{M}[\dots, \mathbf{V}](\overrightarrow{\mathbf{e}_x}) \rightarrow [\mathbf{new} \mathbf{t}_{\mathbf{this}}/\mathbf{this}, \mathbf{new} \mathbf{t}_x/\mathbf{x}] e_y}$$

$$\frac{\mathbf{A}[\mathbf{W}] \triangleleft \mathbf{t}_s \{ \dots \vec{m} \} \in \vec{c} \quad \mathbf{M}[\mathbf{V}](\vec{x}) \{ e \} \in \vec{m}}{\vec{c} \vdash_c [\mathbf{A}[\mathbf{W}].\mathbf{M}[\mathbf{V}]] = (\vec{x}) \{ e \}} \quad \frac{\mathbf{A}[\mathbf{W}] \triangleleft \mathbf{t}_s \{ \dots \vec{m} \} \in \vec{c} \quad \mathbf{M}[\mathbf{V}](\dots) \{ \dots \} \notin \vec{m}}{\vec{c} \vdash_c [\mathbf{t}_s.\mathbf{M}[\mathbf{V}]] = (\vec{x}) \{ e \}} \quad \frac{}{\vec{c} \vdash_c [\mathbf{A}[\mathbf{W}].\mathbf{M}[\mathbf{V}]] = (\vec{x}) \{ e \}}$$

The method lookup operator $[\mathbf{A}[\mathbf{U}].\mathbf{M}[\mathbf{V}]] = \overrightarrow{\mathbf{x} \{ e \}}$ implements basic evaluation virtual method lookup in an OO language, finding the most overriding implementation of method $\mathbf{M}[\mathbf{V}]$ from a given evaluation type $\mathbf{A}[\mathbf{U}]$. Soundness can be stated with the addition of the rule GLOBAL-OK defined in Section 3.1:

$$\frac{\text{Main}[o] \triangleleft \text{Object} \{ \text{Object main}() \} \{ \text{Object main}[\text{Main}[o], o]() \} \{ e_m \} \in \vec{c}}{\vdash_{tc} \mathbf{GLOBAL-OK} \vec{c} \Rightarrow \vec{c} \vdash_e [\mathbf{new} \text{Main}[o]/\mathbf{this}] e_m \rightarrow e_n} \quad \exists \mathbf{A}, \mathbf{V} (e_n = \mathbf{new} \mathbf{A}[\mathbf{V}] \quad \text{or} \quad e_n = \mathbf{downcast_error} \quad \text{or} \quad e_n \text{ is an infinite loop})$$

While core language type checking rules utilize link offsets, these offsets are empty when the rule MODULAR-OK is used to perform type checking. We specify type checking rules that examine linking offsets so GLOBAL-OK can be used to compare the type correctness of the resulting link-reduced classes. Link reduction in MiniJiazzi binds every class reference and method reference expressed in a class expressions to an appropriate linking offset according to the definition of $\vec{s}, u \vdash_l \rightarrow$ given in Section 3.4. As such, link reduction is implemented in the following judgments as a traversal through the implementation of a unit expression:

MiniJiazzi link reduction

$$\begin{array}{c}
\frac{\vec{s}, u \vdash_l \overrightarrow{s_a} \rightarrow \overrightarrow{s_b} \quad \vec{s}, u \vdash_l \overrightarrow{s_a} \mid \overrightarrow{m_a} \rightarrow \overrightarrow{m_b},}{\vec{\forall} \vdash_l u = \dots \{ \overrightarrow{s_a} \mid \overrightarrow{m_a} \} \rightarrow \overrightarrow{s_b} \mid \overrightarrow{m_b} \} \text{ (u-link)}}
\end{array}
\qquad
\frac{\vec{s}, u \vdash_l A[o] \overrightarrow{n_a} \rightarrow \overrightarrow{n_b} \quad \vec{s}, u \vdash_l t_a \rightarrow t_b}{\vec{s}, u \vdash_l A \mapsto W}$$

$$\frac{\vec{s}, u \vdash_l A[o] t_r M[o](t_x) \rightarrow t_v M[W](t_y) \quad \vec{s}, u \vdash_l e_a \rightarrow e_b}{\vec{s}, u \vdash_l A[o] t_r M[o](\overrightarrow{t_x} \mid \overrightarrow{x}) \{ e_a \} \rightarrow t_v M[W](\overrightarrow{t_y} \mid \overrightarrow{x}) \{ e_b \} \text{ (m-link)}}$$

$$\frac{\vec{s}, u \vdash_l t_r \rightarrow t_v \quad \vec{s}, u \vdash_l \overrightarrow{t_x} \rightarrow \overrightarrow{t_y}}{\vec{s}, u \vdash_l A.M \mapsto W}$$

$$\frac{\vec{s}, u \vdash_l A[o] t_r M[o](\overrightarrow{t_x})}{\vec{s}, u \vdash_l A[o] t_r M[W](\overrightarrow{t_x}) \text{ (n-link)}}$$

$$\frac{\vec{s}, u \vdash_l A \mapsto W}{\vec{s}, u \vdash_l A[o] \rightarrow A[W] \text{ (t-link)}}$$

$$\frac{\vec{s}, u \vdash_l t_a \rightarrow t_b}{\vec{s}, u \vdash_l \text{new } t_a \rightarrow \text{new } t_b \text{ (e-link)}}$$

$$\frac{\vec{s}, u \vdash_l t_a \rightarrow t_b \quad \vec{s}, u \vdash_l e_a \rightarrow e_b}{\vec{s}, u \vdash_l (t_a) e_a \rightarrow (t_b) e_b \text{ (e-link)}}$$

$$\frac{\vec{s}, u \vdash_l e_a \rightarrow e_b \quad \vec{s}, u \vdash_l \overrightarrow{e_x} \rightarrow \overrightarrow{e_y} \quad \vec{s}, u \vdash_l B.M \mapsto W \quad \vec{s}, u \vdash_l B \mapsto V}{\vec{s}, u \vdash_l e_a.M[B[o], o](\overrightarrow{e_x}) \rightarrow e_b.M[B[V], W](\overrightarrow{e_y}) \text{ (e-link)}}$$

B Proof of Lemma 4

In MiniJiazzi, the sink and source typing environments used in a congruence check always overlap. The compilation typing environment and the intermediate typing environment of a unit are both formed using the unit's imported class signatures, and the intermediate typing environment of a unit and the program's linkage typing environment are both formed using the unit's exported class signatures. A consequence of source/sink typing environments that express equivalent subtyping relationships is that we can reason about the uncommon parts of the source and sink typing environments with Lemma 5:

Lemma 5 (PARTIAL-SUBTYPING) *For source and sink typing environments that have equivalent subtyping relationships with respect to sink classes, the uncommon parts of the source and sink typing environments, express equivalent subtyping relationships from classes the sink typing environment.*

$$\frac{\overrightarrow{s_{src}} \supseteq \overrightarrow{s_{snk}} \mid \vdash \text{UNIQUE} \overrightarrow{s_{cmn}} \mid \cup \overrightarrow{s_{snk}} \mid \vdash \text{UNIQUE} \overrightarrow{s_{cmn}} \mid \cup \overrightarrow{s_{src}} \mid}{\forall A[V] \in \overrightarrow{s_{snk}} \mid \cup \overrightarrow{s_{cmn}} \mid, \forall B[W] \in \overrightarrow{s_{snk}} \mid \cup \overrightarrow{s_{cmn}} \mid. \overrightarrow{s_{snk}} \cup \overrightarrow{s_{cmn}} \vdash_t A[V] < B[W] \leftrightarrow \overrightarrow{s_{src}} \cup \overrightarrow{s_{cmn}} \vdash_t A[V] < B[W]}$$

$$\frac{\overrightarrow{s_{src}} \supseteq \overrightarrow{s_{snk}} \mid \vdash \text{UNIQUE} \overrightarrow{s_{cmn}} \mid \cup \overrightarrow{s_{snk}} \mid \vdash \text{UNIQUE} \overrightarrow{s_{cmn}} \mid \cup \overrightarrow{s_{src}} \mid}{\forall A[V] \in \overrightarrow{s_{snk}} \mid, \forall B[W] \in \overrightarrow{s_{cmn}} \mid \cup \overrightarrow{s_{snk}} \mid. \overrightarrow{s_{src}} \vdash_t A[V] < B[W] \leftrightarrow \overrightarrow{s_{snk}} \vdash_t A[V] < B[W]}$$

We can take a partial typing environment equivalence relationship and combine them with other partial typing environment equivalent relationships as long as the class signatures involved do not overlap, or if the class signatures overlap completely. We can express this as Lemma 6 and Lemma 7:

Lemma 6 (COMBINE-PARTIAL-1) *Two partial typing environments subtyping equivalence relationships with dis-*

joint sink class signatures can be combined.

$$\frac{\begin{array}{l} \forall A[U] \in \overrightarrow{A[U]} \quad \forall B[V] \in \overrightarrow{B[V]} . \overrightarrow{s}_a \vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{s}_b \vdash_t A[U] < B[V] \quad \vdash \text{UNIQUE} \overrightarrow{s}_a \mid \cup \mid \overrightarrow{s}_c \\ \forall C[W] \in \overrightarrow{C[W]} \quad \forall D[Z] \in \overrightarrow{D[Z]} . \overrightarrow{s}_c \vdash_t C[o] < D[o] \leftrightarrow \overrightarrow{s}_d \vdash_t C[W] < D[Z] \quad \vdash \text{UNIQUE} \overrightarrow{s}_b \mid \cup \mid \overrightarrow{s}_d \\ \overrightarrow{A[o]} \subseteq \mid \overrightarrow{s}_a \mid \quad \overrightarrow{A[U]} \subseteq \mid \overrightarrow{s}_b \mid \quad \overrightarrow{C[o]} \subseteq \mid \overrightarrow{s}_c \mid \quad \overrightarrow{C[W]} \subseteq \mid \overrightarrow{s}_d \mid \end{array}}{\forall A[U] \in \overrightarrow{A[U]} \cup \overrightarrow{C[W]} \quad \forall B[V] \in \overrightarrow{B[V]} \cup \overrightarrow{D[Z]} . \overrightarrow{s}_a \cup \overrightarrow{s}_c \vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{s}_b \cup \overrightarrow{s}_d \vdash_t A[U] < B[V]}$$

Lemma 7 (COMBINE-PARTIAL-2) A partial typing environment subtyping equivalence relationship can be combined with another partial typing environment subtyping equivalence relationship that contains all of its sink class signatures.

$$\frac{\begin{array}{l} \forall A[U] \in \overrightarrow{A[U]} \quad \forall B[V] \in \overrightarrow{B[V]} . \overrightarrow{s}_a \vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{s}_b \vdash_t A[U] < B[V] \quad \vdash \text{UNIQUE} \overrightarrow{s}_a \mid \cup \mid \overrightarrow{s}_c \\ \forall C[W] \in \overrightarrow{C[W]} \quad \forall D[Z] \in \overrightarrow{D[Z]} . \overrightarrow{s}_a \cup \overrightarrow{s}_c \vdash_t C[o] < D[o] \leftrightarrow \overrightarrow{s}_b \cup \overrightarrow{s}_d \vdash_t C[W] < D[Z] \quad \vdash \text{UNIQUE} \overrightarrow{s}_b \mid \cup \mid \overrightarrow{s}_d \\ \overrightarrow{A[o]} \subseteq \mid \overrightarrow{s}_a \mid \quad \overrightarrow{A[U]} \subseteq \mid \overrightarrow{s}_b \mid \quad \overrightarrow{C[o]} \subseteq \mid \overrightarrow{s}_a \mid \cup \mid \overrightarrow{s}_c \mid \quad \overrightarrow{C[W]} \subseteq \mid \overrightarrow{s}_b \mid \cup \mid \overrightarrow{s}_d \mid \end{array}}{\forall A[U] \in \overrightarrow{A[U]} \cup \overrightarrow{C[W]} \quad \forall B[V] \in \overrightarrow{B[V]} \cup \overrightarrow{D[Z]} . \overrightarrow{s}_a \cup \overrightarrow{s}_c \vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{s}_b \cup \overrightarrow{s}_d \vdash_t A[U] < B[V]}$$

Proof of Lemma 4: Given a type correct linked program, we can use Lemma 3 and transitivity that a linked unit's compilation typing environment and the program's linkage typing have equivalent subclassing relationships with respect to all classes in unit intermediate typing environments:

$$\frac{\forall A[o] \in \mid \overrightarrow{s}_{iu} \mid \cup \mid \overrightarrow{s}_{eu} \mid, \forall B[o] \in \mid \overrightarrow{s}_{iu} \mid \cup \mid \overrightarrow{s}_{eu} \mid . \cup \{ \overrightarrow{s}_{iu} \overrightarrow{s}_{eu} \} \{ \overrightarrow{c}_{xu} \} \in \overrightarrow{u}}{\overrightarrow{s}_{iu} \cup \overrightarrow{c}_{xu} \vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{s}_e \vdash_t A[o] < B[o]}$$

Use Lemma 5 to show that the subclassing relationships expressed by a unit's exported class signatures are also expressed in the unit's implementation class expressions:

$$\frac{\forall A[o] \in \mid \overrightarrow{s}_{eu} \mid, \forall B[o] \in \mid \overrightarrow{s}_{iu} \mid \cup \mid \overrightarrow{s}_{eu} \mid . \cup \{ \overrightarrow{s}_{iu} \overrightarrow{s}_{eu} \} \{ \overrightarrow{c}_{xu} \} \in \overrightarrow{u}}{\mid \overrightarrow{c}_{xu} \mid \vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{s}_{eu} \vdash_t A[o] < B[o]}$$

Relate to the linked reduced versions of a unit's implementation classes:

$$\frac{\forall U \in \overrightarrow{V}, \forall A[o] \in \mid \overrightarrow{s}_{eu} \mid, \forall B[o] \in \mid \overrightarrow{s}_{iu} \mid \cup \mid \overrightarrow{s}_{eu} \mid . u = \cup \{ \overrightarrow{s}_{iu} \overrightarrow{s}_{eu} \} \{ \overrightarrow{c}_{xu} \} \in \overrightarrow{u}}{\overrightarrow{s}, u \vdash_l B \mapsto W \quad \overrightarrow{s}, u \vdash_l \overrightarrow{c}_{xu} \rightarrow \overrightarrow{c}_{yu} \quad \overrightarrow{c}_{yu} \vdash_t A[U] < B[W] \leftrightarrow \overrightarrow{s}_{eu} \vdash_t A[o] < B[o]}$$

Use Lemma 6 to recombine the partial typing environments equivalent subclassing relationships into complete typing environment equivalent subclassing relationships:

$$\forall A[o] \in \mid \overrightarrow{s}_{eu} \mid, \forall B[o] \in \mid \overrightarrow{s}_{eu} \mid . \cup \overrightarrow{s}_{eu} \in \overrightarrow{V} \overrightarrow{s}_e \quad W \overrightarrow{s}_{eu} \in \overrightarrow{V} \overrightarrow{s}_e \quad \mid \overrightarrow{c}_y \mid \vdash_t A[U] < B[W] \leftrightarrow \overrightarrow{s}_e \vdash_t A[o] < B[o]$$

Consider each unit one at a time, and only its imported and exported classes, from the previous step we can derive the following:

$$\forall U \in \vec{V}, \forall A[o] \in \overrightarrow{\mathfrak{s}_{iu}} \cup \overrightarrow{\mathfrak{s}_{eu}}, \forall B[o] \in \overrightarrow{\mathfrak{s}_{iu}} \cup \overrightarrow{\mathfrak{s}_{eu}}. \quad \mathbf{u} = U \{ \overrightarrow{\mathfrak{s}_{iu}} \overrightarrow{\mathfrak{s}_{eu}} \} \{ \overrightarrow{\mathfrak{c}_{xu}} \} \in \vec{\mathbf{u}}$$

$$\vec{\mathfrak{s}}, \mathbf{u} \vdash_l A \mapsto V \quad \vec{\mathfrak{s}}, \mathbf{u} \vdash_l B \mapsto W \quad \overrightarrow{\mathfrak{s}_{iu}} \cup \overrightarrow{\mathfrak{c}_{xu}} \Vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{\mathfrak{c}_y} \Vdash_t A[V] < B[W]$$

We can use transitivity to relate equivalent subclassing relationships between the link-reduced typing environment and compilation typing environments, but only for classes in the intermediate typing environment:

$$\forall U \in \vec{V}, \forall A[o] \in \overrightarrow{\|\mathfrak{c}_{xu}\|}, \forall B[o] \in \overrightarrow{\mathfrak{s}_{iu}} \cup \overrightarrow{\|\mathfrak{c}_{xu}\|}. \quad \mathbf{u} = U \{ \overrightarrow{\mathfrak{s}_{iu}} \overrightarrow{\mathfrak{s}_{eu}} \} \{ \overrightarrow{\mathfrak{c}_{xu}} \} \in \vec{\mathbf{u}}$$

$$\vec{\mathfrak{s}}, \mathbf{u} \vdash_l B \mapsto W \quad \vec{\mathfrak{s}}, \mathbf{u} \vdash_l \overrightarrow{\mathfrak{c}_{xu}} \rightarrow \overrightarrow{\mathfrak{c}_{yu}} \quad \overrightarrow{\|\mathfrak{c}_{xu}\|} \Vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{\mathfrak{c}_{yu}} \Vdash_t A[U] < B[W]$$

And we can use Lemma 7 to combine this result with the previous result to complete the proof of Lemma 4:

$$\forall U \in \vec{V}, \forall A[o] \in \overrightarrow{\mathfrak{s}_{iu}} \cup \overrightarrow{\|\mathfrak{c}_{xu}\|}, \forall B[o] \in \overrightarrow{\mathfrak{s}_{iu}} \cup \overrightarrow{\|\mathfrak{c}_{xu}\|}. \quad \mathbf{u} = U \{ \overrightarrow{\mathfrak{s}_{iu}} \overrightarrow{\mathfrak{s}_{eu}} \} \{ \overrightarrow{\mathfrak{c}_{xu}} \} \in \vec{\mathbf{u}}$$

$$\vec{\mathfrak{s}}, \mathbf{u} \vdash_l B \mapsto W \quad \overrightarrow{\mathfrak{s}_{iu}} \cup \overrightarrow{\|\mathfrak{c}_{xu}\|} \Vdash_t A[o] < B[o] \leftrightarrow \overrightarrow{\mathfrak{c}_y} \Vdash_t A[U] < B[W]$$