

cJ: Enhancing Java with Safe Type Conditions

Shan Shan Huang David Zook

College of Computing
Georgia Institute of Technology
ssh,dzook@cc.gatech.edu

Yannis Smaragdakis

Department of Computer and Information Science
University of Oregon
yannis@cs.uoregon.edu

Abstract

cJ is an extension of Java that allows supertypes, fields, and methods of a class or interface to be provided only under some static subtyping condition. For instance, a cJ generic class, $C\langle P \rangle$, may provide a member method m only when the type provided for parameter P is a subtype of a specific type Q .

From a practical standpoint, cJ adds to generic Java classes and interfaces the ability to express case-specific code. Unlike conditional compilation techniques (e.g., the C/C++ “`#ifdef`” construct) cJ is statically type safe and maintains the modular type-checking properties of Java generic classes: a cJ generic class can be checked independently of the code that uses it. Just like regular Java, checking a cJ class implies that all uses are safe, under the contract for type parameters specified in the class’s signature.

As a specific application, cJ addresses the well-known shortcomings of the Java Collections Framework (JCF). JCF data structures often throw run-time errors when an “optional” method is called upon an object that does not support it. Within the constraints of standard Java, the authors of the JCF had to either sacrifice static type safety or suffer a combinatorial explosion of the number of types involved. cJ avoids both problems, maintaining both static safety and conciseness.

Categories and Subject Descriptors D.1.2 [Programming Techniques]: Automatic Programming—program synthesis, program transformation, program verification; D.1.5 [Programming Techniques]: Object-oriented Programming; D.2.13 [Software Engineering]: Reusable Software—Reusable libraries; D.3.3 [Programming Languages]: Language Constructs and Features; D.3.4 [Programming Languages]: Processors

General Terms Design, Languages

Keywords aspect-oriented programming, meta-programming, language extensions

1. Introduction

Generic types increase the expressiveness and safety of a programming language. Since the introduction of Java and C#, researchers have worked on adding genericity mechanisms that were subsequently integrated into the base languages themselves [3, 15, 31]. From a language design standpoint, modern genericity mechanisms

offer a good tradeoff between expressiveness and separate checkability. For instance, Java generics have limited expressiveness compared to undisciplined mechanisms, such as C++ template classes, yet offer the ability to detect static errors (e.g., type errors) without having to provide a specific type parameter that triggers the error.

This paper proposes cJ: an extension of Java that adds more expressiveness to its genericity mechanism without sacrificing any of the Java type-checking guarantees. Specifically, we add to Java the ability to place type-conditions on methods, fields, or supertypes. This is best illustrated with a small example. Consider the following generic cJ class:

```
class C<X> {
  X xRef;
  ...
  <X extends DataSource?>
  void store() { ... xRef.getConnection() ... }
}
```

In this example, the member method `store` is declared in a type-instantiation of generic class `C`, *only* when the type argument for `X` is a class (or interface) that implements (resp., extends) interface `DataSource`. The `<...?>` syntax is cJ’s type-conditional construct. One can read this syntax as “static-if”, or just “if”. The call `xRef.getConnection()` is well-typed only because type `X` is guaranteed to be a subtype of `DataSource` and, consequently, to provide the `getConnection` method.

cJ is translated by erasure, reducing to regular Java in a backward compatible manner. This allows us to solve a well-recognized problem in the Java Collections Framework (JCF), the standard Java data structures library. Currently, JCF data structures support two main common interfaces (`Collection` and `Map`), regardless of optional behavior, such as whether the data structure is modifiable or not, and whether the data structure has variable size or not. Classes that do not support the corresponding operations throw `UnsupportedOperationException` when the operations are called at run-time. The design of the JCF is an instance of sacrificing static type safety in favor of conciseness. cJ solves this problem, maintaining both type safety and conciseness of expression.

Interestingly, cJ can be thought of as a language that allows cross-cutting [17] at the level of types. cJ type conditions are used to define many implicit types from a single class definition. In the above example, the single definition of class `C` can be thought of as defining the implicit types `C<subtype-of-DataSource>` and `C<not-DataSource>`. Thus, with cJ type conditionals, one can add orthogonal “aspects” or “dimensions” to an existing type hierarchy. Our re-implementation of the JCF provides a vivid demonstration of this feature. Beginning from a simple subtyping hierarchy, we introduce variation based on whether a data structure supports content modification or size variability. The definitions of the various collections (e.g., the `List`, `Collection`, `Map`, and `Set` interfaces) are as simple as in plain Java, yet a much richer type hierarchy is produced by modifying each type with attributes chosen

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AOSD 07 March 12-16, 2007, Vancouver, BC, Canada
Copyright © 2007 ACM 1-59593-615-7/07/03...\$5.00

from a separate type hierarchy (with types such as `Modifiable`, `DeleteOnly`, and `Resizable`). Thus, whereas traditional AOP tools allow the expression of cross-cutting features at the level of methods, cJ supports separation of concerns at the type level. Type hierarchies can be specified separately to represent orthogonal concerns, and cJ allows their composition to form richer, derived hierarchies. Writing general code that exploits these derived hierarchies can be done through a natural extension of the Java variance/wildcards mechanism (e.g., “`? extends T`” clauses).

Although Java was chosen as the platform for our ideas, the cJ approach is far from Java-specific. The same programming and type-checking framework can be applied to other languages. Yet Java is a good representative of modern OO languages and integrating with it demonstrates clearly both the benefits and the intricacies of our approach.

Concretely, our work makes the following contributions:

- cJ allows expressing highly variable generic classes concisely. Compared to standard OO mechanisms, cJ allows a single generic class to express the functionality of an exponential number of regular Java classes.
- cJ offers full type safety, analogous to that of the base Java language. A cJ generic class is checked separately from its uses. The type system ensures that the class is type-correct under any consistent combination of outcomes of the type-conditionals.
- Other research work [7, 19, 20, 22] has targeted the problem of type-safe conditional declarations. Nevertheless, many of the past mechanisms are in a simpler context (e.g., no subtyping) or do not allow some of the cJ features (e.g., conditional subtyping). None of the past research dealt with the integration of conditional members and subtypes with (use-site) variance, or provided a backward compatible, erasure-based translation. Overall, cJ is distinguished by its power and its smooth integration in a modern language.
- cJ solves the static type safety issues of the JCF. We are not aware of other language proposals that address this well-publicized need without sacrificing conciseness.

The rest of the paper is organized as follows. We first give an informal introduction to the cJ language extensions. This serves as background for the motivating examples of Section 3 and the JCF case study. In Section 4, we present interesting ways in which the cJ extensions interact with variance in Java. We then analyze the cJ implementation in Section 5. Section 6 formalizes cJ and subsequently we discuss related work (Section 7) and our conclusions.

2. cJ Language Introduction

We next give an informal overview of cJ’s syntax and semantics, to prepare the ground for our motivating examples. A formal description of the language is laid out in Section 6.

2.1 cJ Basics and Examples

cJ is a conservative extension of Java—we assume Java 5, with support for generics and variance (“wildcards”) [3, 29] in the base language. cJ adds to Java the ability to change a type’s structure depending on static type conditions. The language provides a *type-conditional* construct. The following is a simple example showing the use of a type-conditional:

```
class Foo<T> {
  <T extends Bar?>
  int i;
}
```

In the above example, `<T extends Bar?>` is a type-conditional. The declaration immediately following it, “`int i;`”, exists only if `Foo` is parameterized by a subtype of `Bar`.

Type-conditionals can be used for the declaration of class- or interface-level methods and fields, as well as for the declaration of conditional supertypes. For instance, we can have:

```
class Foo<T>
<T extends Serializable?> implements Serializable
{ ... }
```

The above class `Foo` implements interface `Serializable` only when it is parameterized by a type that also implements (or extends) `Serializable`.

Multiple declarations can exist in the same conditional block by surrounding them in `< ... >`. For instance:

```
class Foo<T> {
  <T extends Bar?>
  < int i;
    void meth(T t) { }
  >
}
```

The above is equivalent to preceding each declaration individually with the type-conditional `<T extends Bar?>`.

There are two required components to each type-conditional block. The first is the type condition, defined inside “`< ... >?`”. Any syntax that is valid for defining the type parameters of a Java class is valid here as a type condition: the type conditions have the standard F-bounded polymorphism form [3], where a type parameter can be referenced by its own bound, e.g., “`T extends I<T>`”. Note especially that “`extends`” is used to express all kinds of subtyping constraints (including interface conformance) and that the syntax admits conjunctions of subtyping bounds (e.g., “`T extends I<T> & J<T>`”), as well as bounding multiple parameters (e.g., “`S extends I<S>, T extends J<S>`”). Similarly to Java, it is *not* valid for a type parameter to appear by itself on the right hand side of `extends` (i.e. we cannot place lower bounds on a type parameter). Also note that only type parameters declared at the class/interface level are allowed in type conditions. Polymorphic method type parameters are not allowed. The second required component, the consequent block, immediately follows the type condition. Declarations within this block exist for the enclosing type if and only if the type condition is true, after all type parameters are instantiated. A type-conditioned declaration is syntactically a declaration, hence, type-conditionals can nest.

cJ ensures that all uses of type-conditionals are statically safe. All code should be well-typed under its enclosing type conditions. Furthermore, all uses of class or interface members should be under equivalent or stronger conditions than those employed in the member’s declaration. For instance, the following use is legal only if `J` is a subtype of `I`:

```
class Foo<T> {
  <T extends I?>
  int i;

  <T extends J?>
  void incI() { i++; } // legal iff J subtypes I
}
```

The following code is also legal, as the type conditions are strengthened by adding conjunctions:

```
class Foo<T,U> {
  <T extends I?>
  int i;

  <T extends I, U extends K?>
  void incI() { i++; } // legal: stronger condition

  <T extends I & J?>
  void decI() { i--; } // legal: stronger condition
}
```

2.2 Restrictions

There are some restrictions that cJ imposes on conditional declarations. These restrictions significantly simplify the translation and interfacing with existing Java code, as we will discuss in Section 5. The rule of thumb is that a cJ class (or interface) should be a legal Java class (interface) if all type conditions are removed.

A cJ class can have at most one `extends` clause, regardless of whether it is under a type-conditional. Of course, a cJ class can implement multiple interfaces and any of the `implements` clauses can be conditional.

Declarations that are conflicting per the standard Java rules are not allowed, even if their type-conditional conditions are exclusive. For instance, the following is illegal in cJ, even when neither `Baz` nor `Bar` are a subtype of the other:

```
interface IFoo<T> {
    <T extends Bar>?
    void foo(int i);

    <T extends Baz>?
    int foo(int i); // duplicate definition
}
```

Furthermore, subtypes are required to define conditional methods under equivalent or weaker conditions than conflicting methods in their (possibly conditional) supertypes. For example:

```
interface ISuper<T,U> {
    <T extends Bar>?
    void meth1(int i);

    <T extends Bar>?
    void meth2(Object o);
}

interface ISub<T,U> extends ISuper<T,U> {
    <T extends Bar & Baz>?
    void meth1(int i); // illegal unless Bar subtypes Baz

    void meth2(Object o); // legal: weaker (no condition)
}
```

The above rules extend to the members of conditional supertypes. Their type conditions from the perspective of the subtype is the conjunction of the subtyping and the membership conditions. (E.g., a member under a static condition P in an interface implemented under condition Q should be thought of as being under a condition $P \& Q$ for the purposes of the above discussion.) Our formalism in Section 6 makes this definition precise.

3. cJ Benefits

Having introduced the cJ language, we can now examine some motivating examples. We first discuss a small example that demonstrates how a type-conditional avoids a combinatorial blowup of the number of classes required in a Java application. Then, we examine a specific case study: the Java Collections Framework and its well-known shortcomings with respect to static type safety.

3.1 The Argument for Safety and Conciseness

There are two ways to view the benefits of cJ over regular Java. In Java, when the contents of a class can vary with respect to multiple orthogonal concerns, the programmer can either choose to maintain static type safety and suffer a combinatorial explosion of the number of classes involved, or sacrifice static type safety in order to keep the number of classes manageable. cJ achieves both benefits simultaneously.

The conciseness benefits of cJ are relatively easy to see. When multiple conditionals capture different axes of variability, a cJ generic class corresponds to a hierarchy of many different regular classes. Consider a simple example class `C`:

```
class C<X> {
    <X extends Serializable>?
    public void store() { ... }
    ...
    <X extends Comparable<X>>?
    public X getMin() { ... }
}
```

That is, class `C` supports method `store` only when type parameter `X` is a serializable type. Similarly, `C` supports method `getMin` only when type parameter `X` is a comparable type.

To achieve the same effect with regular Java, the programmer needs to create separate classes that capture all possible combinations. One possibility would be the following class hierarchy:

```
class CommonC<X> {
    ... // the common parts of C
}

class CSer<X extends Serializable> extends CommonC<X> {
    public void store() { ... }
}

class CComp<X extends Comparable<X>> extends CommonC<X> {
    public X getMin() { ... }
}

class CCompSer<X extends Comparable<X> & Serializable>
    extends CSer<X>
{
    public X getMin() { ... }
}
```

The result is four different classes, capturing the same content as the original cJ class. Method code is replicated: `CCompSer` cannot inherit `getMin` from `CComp` because it already has a superclass, `CSer`. Furthermore, `CCompSer` is not a subtype of `CComp`, hence, a `CCompSer` object cannot be used where a `CComp` is expected, even though it supports the required methods of a `CComp`. Such code replication and subtyping problems can be alleviated by using delegation techniques and interfaces, but this may require significant code reorganization, weakening of encapsulation, and explicitly maintaining object identity. For instance, to minimize code length with delegation, the programmer often needs to enable access to members of another class, as well as manually ensure a one-to-one mapping among different sub-objects.

Note that this example deals with only two axes of variability: whether `X` is `Comparable` and whether `X` is `Serializable`. Still, the result is undesirable. In the general case, the number of Java classes required for a faithful emulation is exponential to the number of distinct type-conditionals in the cJ class, assuming a straightforward mapping. Overall code length will also be exponentially greater, unless delegation, with its aforementioned disadvantages, is used.

In practice, it is unlikely that Java developers would want to deal with this kind of combinatorial complexity. Instead, they will likely prefer to provide a single type that captures the union of all possible members. In that case, when an “unsupported” method is called, a run-time error can be signaled in the form of an exception. For instance, following Java conventions, our earlier example is likely to be written in standard Java as follows:

```
class C<X> {
    public void store()
        throws UnsupportedOperationException
    { ... }
    ...
    public X getMin()
        throws UnsupportedOperationException
    { ... }
}
```

This addresses the code size and number-of-types explosion problem at the expense of sacrificing static type safety. The type checker is no longer able to tell under what conditions the `store` and `getMin` operations would be illegal. A run-time type error is produced instead, when illegal operations get called.¹ Relative to plain Java, cJ combines the advantages of static type safety and code conciseness.

It is worth noting that the cJ compiler translates its input into plain Java by following an approach similar to that of the example above (i.e., a single class is produced, containing all possible members). Yet, the cJ type system statically ensures that no exceptions for unsupported methods are thrown at run-time. We describe the cJ implementation in Section 5.

Finally, an interesting question on the power of type-conditionals concerns their value under multiple inheritance. Multiple inheritance can address the problems of delegation, in that it allows composing a class modularly without violating object identity or encapsulation. If Java had multiple inheritance, in addition to its bounded generics, the above example could be expressed in the same amount of code as in cJ. Nevertheless, the main benefit of type-conditionals is not in minimizing the code length, but in minimizing the number of explicit types that users need to manage. Consider the case of a type hierarchy among types `I1`, `I2`, ..., `IN`. Type conditionals allow the programmer to create implicitly a virtual isomorphic hierarchy by using a single class `C<X>` with member and/or supertype declarations conditional on `X` extending `I1`, `I2`, etc. The language will automatically ensure that the two hierarchies have consistent structure. If, for instance, `I1` is a subtype of `I5`, all methods in `C` declared conditionally under `X extends I1` will be able to access methods declared conditionally under `X extends I5`. With traditional subtyping mechanisms, the user would need to create explicit types `C1`, `C2`, ..., `CN` with a subtyping hierarchy reflecting the one of `I1`, `I2`, ..., `IN`. Relieving the programmer from explicitly managing these types is the greatest advantage of type-conditionals in *any* language setting. As we discuss in the next section, the stated motivation of Java developers for choosing a type-unsafe solution for the JCF was not avoiding code size explosion but avoiding an explosion in the number of explicit types that users would need to deal with.

3.2 Case Study: Java Collections Framework

A striking demonstration of the problems presented above can be found in the Java Collections Framework: the standard Java data structures library. The JCF supplies types such as `Collection`, `Set`, `Map`, and `List`. However, there are other cross-cutting concerns along which to organize these basic data structures. One such concern is that of “modifiability”: is a data structure modifiable through its public interface or not? This concept is not captured via the Java type system in the design of the JCF. Instead, any attempt to modify an “unmodifiable” collection results in the throwing of an `UnsupportedOperationException` at run-time. Another similar concern is that of size variability. Some data structures are modifiable, yet their size cannot change—arrays are a standard example. An array supports the operations of the `List` interface with the exception of `add` or `remove`, which throw `UnsupportedOperationException`. This is a case of circumventing the static type system in order to avoid a combinatorial explosion in the number of types specified in the library. In fact, six out of the fifteen methods of interface `Collection` in JDK 1.5 are optional and may result in run-time errors.

¹The `UnsupportedOperationException` is a run-time exception (i.e., the compiler does not check that it is always caught or declared) and a member of the JCF. For the purposes of this paper, we use this exception type even for code outside the JCF. Any different exception could assume the same general role.

The above is a well-known issue. The very first “frequently asked question” in the Java Collections API Design FAQ² is:

Why don't you support immutability directly in the core collection interfaces so that you can do away with optional operations (and `UnsupportedOperationException`)?

The design rationale reflected in the answer to this FAQ indirectly offers a compelling argument for cJ. The developers note:

Clearly, static (compile time) type checking is highly desirable, and is the norm in Java. We would have supported it if we believed it were feasible. Unfortunately, attempts to achieve this goal cause an explosion in the size of the interface hierarchy ...

Subsequently, the Java Collections API developers proceed to give an illustration of the kinds of “explosion in size” problems that a type-safe design would encounter, if cross-cutting concerns such as “modifiable”, “variable-size”, “append-only”, etc., are expressed in the type system. The Java Collections Design FAQ concludes:

Now we're up to twenty or so interfaces and five iterators, and it is almost certain that there are still collections arising in practice that don't fit cleanly into any of the interfaces.

The above issue is not specific to the Java Collections Framework. Other developers of Java data structure libraries have identified the same shortcomings. Doug Lea (quoted in the JCF FAQ) authored a popular Java collections package and remarks:

Much as it pains me to say it, strong static typing does not work for collection interfaces in Java.

(We invite the reader to consult online the informative FAQ answer, which we cannot reproduce here in its entirety.)

cJ addresses fully and cleanly the above problem with the JCF. Interfaces `Collection`, `List`, etc. are implemented modularly using type-conditionals. Specifically, there are three interesting properties that we capture: whether a collection is modifiable, whether it supports only deletions, and whether it supports both deletions and additions (i.e., all size change operations). These cross-cutting concerns are expressed using (marker) interfaces `Modifiable`, `DeleteOnly` and `Resizable`. The `Resizable` interface is a subtype of `DeleteOnly`—a resizable collection supports operations such as `clear` and `remove`, but also `add` and `addAll`. By combining these interfaces one can specify different flavors of each collection. This is done through a type parameter `M` passed to each collection generic class. For instance, interfaces `Collection` and `List` are implemented as follows:³

```
interface Collection<E,M> extends Iterable<E,M> {
    < M extends Resizable>?
    <
    boolean add(E o);
    boolean addAll(Collection<? extends E, ?> c);
    >
    <M extends DeleteOnly>?
    <
    boolean removeAll(Collection<?, ?> c);
    void clear();
    ...
    >
    boolean contains(Object o);
    boolean isEmpty();
    ... // other methods common to all collections
}
```

²<http://java.sun.com/j2se/1.5.0/docs/guide/collections/designfaq.html>

³Our re-implementation of the JCF can be found on the cJ website: <http://www.cc.gatech.edu/~ssh/cj>.

```

interface List<E,M> extends Collection<E,M> {
    <M extends Resizable>?
    <
    void add(int index, E element);
    boolean addAll(int index, Collection<? extends E,?> c);
    >
    <M extends DeleteOnly>?
    <
    E remove(int index);
    ...
    >
    <M extends Modifiable>?
    E set(int index, E element);
    ... // other methods common to all lists
}

```

(Note that the question-mark symbol is used both in our type-conditional syntax, and as a wildcard in order to specify variance in generic operations, per the standard Java syntax.) Concrete classes that implement these interfaces (e.g., `ArrayList`) have similarly structured type-conditionals. This implementation is concise without sacrificing static type safety. The user of the `List` interface explicitly selects the desired flavor of the collection. For instance, a possible type instantiation of `List` is `List<Integer,Modifiable>`, signifying a modifiable (but not resizable) list of integers. Another possible instantiation is `List<Integer,Object>` (or any type that is not a subtype of `Modifiable` in place of `Object`) to signify a non-modifiable and non-resizable list. The programmer cannot accidentally call a `set` method on a collection that is statically specified to be unmodifiable. The need for an `UnsupportedOperationException` is eliminated.

The JCF case study serves well as a motivating example for the more powerful cJ features described in later sections. Specifically, the major question we have not yet addressed is how to write general code that abstracts over multiple cJ types. There are two ways to safely abstract over types in the Java type system. One way is to use interfaces—e.g., we may want to write code that works with all `Comparable` objects regardless of whether they are of type `Integer`, `String`, `Array`, etc. The other way is to use variance—e.g., we can write code that works with all `List<X>` objects, as long as the element type, `X`, is a subtype of a given type, say, `Number`. Both of these valuable mechanisms are straightforwardly extended and enhanced in cJ.

cJ conditional supertypes enable abstraction using interfaces even for types that support the corresponding operations only conditionally. For instance, we can have definitions such as:

```

class ArrayList<X,M>
<X extends Comparable<X>>? implements Comparable<List<X>>
{
    <X extends Comparable<X>>?
    public int compareTo(List<X> that) { ... }
    ...
}

```

The above `ArrayList` class implements interface `Comparable` and provides the appropriate `compareTo` method only if its parameter type is also a `Comparable`.⁴ Thus, we can use such `ArrayList` objects with code accepting any `Comparable` object—unlike the original Java `ArrayList` class.

The second kind of abstraction is quite interesting and practically valuable in the cJ setting. For instance, how can we write code that deals uniformly with `List` objects that support at least a `remove` operation, regardless of whether the objects are of type `ArrayList<E,DeleteOnly>` or `ArrayList<E,Resizable>` or any other compatible subtype and “flavor” combination? This is precisely the role of the question-mark wildcard types that appeared

⁴Our thanks to Phil Wadler for this motivating example.

in our above Java Collections code—e.g., for method `addAll`. The general approach follows a natural extension of the standard Java variance mechanism. We discuss this topic in the next section.

4. Subtyping and Variance

cJ type-conditionals turn out to fit very well in the Java type checking framework. In particular, the relationships among different instantiations of the same generic cJ class fall out very simply from the standard rules for variance, with only a small addition. We next give a bird’s eye view of wildcards and variance in the Java type system (readers familiar with variance can skip Section 4.1) and then discuss how these relate to cJ.

4.1 Variance and Wildcards

Here we only give a brief (and simplified) summary of Java wildcards as used to implement variance. A thorough treatment can be found in past literature [12, 29].

Java allows using generic types with a non-specific instantiation, through the wildcard syntax “`? extends T`”, “`? super T`” and “`?`”. For a generic type `C`, the meaning of a `C<? extends T>` is “`C` instantiated with any subtype of `T`”. For instance, the JCF `Collection` interface supports a method:

```

interface Collection<E> extends Iterable<E> {
    ...
    addAll(Collection<? extends E> c);
}

```

The wildcard means that if, for instance, we have an object of type `Collection<Number>`, we can pass as an argument to its `addAll` method an object of type `Collection-of-some-subtype-of-Number`. For instance (assuming `Integer` subtypes `Number`):

```

Collection<Number> c = new ArrayList<Number>();
Collection<Integer> ci = new ArrayList<Integer>();
... // populate ci
c.addAll(ci);

```

Similarly, the wildcard syntax “`C<? super T>`” means “`C` instantiated with any supertype of `T`”, and the syntax “`C<?>`” means “`C` instantiated with anything”.

Wildcards form an elegant way to write highly general code that can apply to multiple instantiations of generic types. Nevertheless, to statically ensure that the result is safe (i.e., that the object can indeed support all the operations that the code wants to perform on it) several restrictions need to be imposed.

- An object `c` of type `C<? extends T>` can only be used to call methods where the type parameter of `C` is in a *co-variant* position, i.e., appears only as the return type of a method, if at all. Also, fields of `c` typed as the type parameter of `C` can only be read from, not written to. For instance, given an object `c` of type `Collection<? extends E>`, we can never invoke a method such as “`boolean add(E o)`” on `c`, because this method is declared in interface `Collection<E>`, and the type parameter `E` appears as an argument type to `add`.
- Similarly, an object `c` of type `C<? super T>` can only be used to call methods with the type parameter of `C` in a *contra-variant* position, i.e., it appears only as an argument type to a method, if at all. Fields of `c` typed as the type parameter of `C` can be written to with values typed `T`, but only read as values of type `Object`.
- An object `c` of type `C<?>` can only be used to call methods where the type parameter of `C` does not appear at all (*bi-variance*). Similarly, the fields of `c` typed as the type parameter of `C` can only be read as `Objects`, and not written to.

Next we discuss how a slight extension of the Java variance rules makes them apply transparently to cJ.

4.2 Variance and Type-Conditionals

We return to the original question regarding type-conditionals and subtyping. Consider a cJ class `C<X>`. Can we write code that is general enough to work type-safely with multiple instantiations of `C<X>` (i.e., for multiple values of `X`)? Consider the simple example from Section 3.1:

```
class C<X> {
  <X extends Serializable?>
  public void store() { ... }
  ...
  <X extends Comparable<X>>?
  public X getMin() { ... }
}
```

Intuitively, `X` is used in this example only in order to add more members to generic class `C`. Thus, a “stronger” `X` (i.e., one that will satisfy more “extends” type conditions) will only result in more members being added. In other words, if type `S` is a subtype of `T` then `C<S>` could safely be a subtype of `C<T>`—generic class `C` can be co-variant in its type parameter.

cJ, just like regular Java, does not automatically relate different instantiations of a generic class via subtyping. That is, in the Java and cJ type systems, an instantiation `C<A>` is never a subtype of `C` for two distinct classes `A` and `B`, regardless of the contents of `C` or how `A` and `B` are related. However, if `A` is a subtype of `B`, then `C<A>` is a subtype of `C<? extends B>` and `C` is a subtype of `C<? super A>`. The programmer can use such subtyping relations to write code that applies to multiple instantiations of a generic class and the language statically checks that the code is safe, based on the rules outlined earlier.

cJ enhances the variance rules to deal with type conditions. For instance, we can have the following method, accepting an argument of the above type `C`:

```
void export(C<? extends Serializable> c) {
  ... c.store(); ...
}
```

That is, the `export` method accepts objects of type `C`-of-some-subtype-of-`Serializable`. The language ensures that the body of `export` uses its argument `c` correctly. In this case, the call to `store` is statically type safe, since for any subtype `X` of `Serializable`, type `C<X>` will support `store`.

The general rule for interactions between type parameters and variance is simple:

An occurrence of type parameter `X` in an `<X extends ...>?` type-condition (on either a supertype declaration or a member declaration) constitutes a co-variant use.

Enhanced with the above rule, all other rules of the standard variance framework of Java apply and enable general type safety.

Consider, for instance, a `Queue` that supports averaging of its elements if they are `Numbers`. (This is an artificial example—the functionality is not part of the Java Collections Framework.):

```
interface Queue<X> {
  <X extends Number?>
  X average();
  ... // other methods
}
```

Both appearances of type parameter `X` are in co-variant positions: either in a type condition, or as a return type. In this case, a method can accept objects of type `Queue`-of-some-subtype-of-`Number` and call `average` on them safely. For instance, we can have:

```
void covariant(Queue<? extends Number> q) {
  Number a = q.average();
}
```

We already saw uses of variance in our Java Collections API case study. Consider the following excerpt from the definition of `Collection<E,M>` in Section 3.2:

```
interface Collection<E,M> extends Iterable<E,M> {
  ...
  boolean addAll(Collection<? extends E, ?> c);
  ...
  boolean removeAll(Collection<?, ?> c);
}
```

Methods `addAll` and `removeAll` in the above use arguments bi-variant with respect to the second type parameter of `Collection`. That is, these methods can accept any collection, regardless of whether it is modifiable or not, delete-only or not, etc. Note that the above type signatures statically prevent the implementation of methods `addAll` and `removeAll` from calling methods such as `add`, `clear`, or `set` on their argument `c`: all these methods are declared conditionally and `c` may not support them. Intuitively, this reflects the intent of the interface for methods `addAll` and `removeAll`: they modify the object from which they are invoked, but not their argument object, from which they only read values to add or remove.

Overall, cJ type-conditionals are an excellent match for Java variance. Not only does variance offer a natural abstraction mechanism for conditional types, but also variance and type-conditionals offer the same kind of benefit in a programming language. Both mechanisms allow specifying a single class `C<X>` and having the type system automatically compute several useful derivative types. In the case of variance the derivative types are `C<? extends T>`, `C<? super T>` and `C<?>`, which contain only the co-variant, contra-variant, and bi-variant methods of the class, with respect to some type `T`. In the case of cJ the derivative types correspond to all possible outcomes of type-conditionals. For instance, `Modifiable-and-DeleteOnly-List` is an implicit type produced from the `List<E,M>` definition. Each cJ implicit type contains only the members that exist for this combination of conditions.

5. Implementation

The design of cJ was carefully planned to admit a simple *erasure-based* translation that is backward compatible with Java code. Each cJ generic class can be translated to a single Java generic class (which in turn can be translated to a single non-generic Java class, per the standard erasure translation of Java generics). This and other implementation topics are discussed next.

Erasure. The current cJ compiler is a source-to-source translator into Java. Nevertheless, exactly the same techniques could be used in a direct-to-bytecode translation. Indeed, the source-to-source translation has even more transparency requirements and demonstrates how well cJ fits the Java model.

cJ translates a class (or interface) with type-conditionals into a Java class (resp., interface) by removing all conditional statements. This enables a single class to play the role of all possible instantiations. Consider our earlier example:

```
class C<X> {
  <X extends Serializable?>
  public void store() { ... }
  ...
  <X extends Comparable<X>>?
  public X getMin() { ... }
}
```

cJ translates `C` into a class:

```
class C<X> {
  public void store() { ... }
  ...
  public X getMin() { ... }
}
```

Note that there is no need for a run-time exception. The cJ type system ensures statically that unsupported methods can never be called. (If client code is not compiled with the cJ compiler, there is no such guarantee. We later discuss how the user can explicitly request dynamic checks to ensure that these methods are not called.)

Erasure Intricacies. The cJ translation requires a few more steps than simply removing the type-conditionals. First, the cJ compiler translates the bodies of conditional methods using type casts that ensure the appropriate type conditions. Second, it supplies dummy method bodies to classes implementing (or extending) an interface (class) with unsupported methods. Lastly, it translates certain type instantiations into their “raw type” forms, and performs the same code generation that a plain Java compiler performs in translating generic code into non-generic code. We demonstrate these translations via examples.

When translating conditional code, the cJ compiler needs to maintain known type bounds for each expression. If this is different from the type the expression would have when conditionals are eliminated, then casts need to be output. Consider the example from the Introduction:

```
class C<X> {
    X xRef;
    ...
    <X extends DataSource>?
    void store() { ... xRef.getConnection() ... }
}
```

The call to `getConnection` is only valid because the type `xRef` is known to be a subtype of `DataSource`. Thus, the compiler needs to emit a cast that will ensure this type constraint when the type-conditional is removed. The cast cannot fail at run-time, as the cJ static type checker ensures the `store` method is only called when `X` is indeed a subtype of `DataSource`. The translated code is:

```
class C<X> {
    X xRef;
    ...
    void store()
    { ...((DataSource) xRef).getConnection()... }
}
```

In the case of interfaces (or abstract classes), our erasure translation means that classes implementing (extending) an interface (abstract class) may need to be automatically enhanced. Consider a conditional interface method. Erasure removes the type-conditional and the method will be declared for all instantiations of the interface. Yet, classes implementing some of these instantiations will not provide implementations of the method, as the method is undeclared for the given type parameters. For instance:

```
interface List<E,M> extends Collection<E,M> {
    <M extends DeleteOnly>?
    E remove(int index);
    ...
}
```

```
class FixedList<E> implements List<E,Object> {
    ... // no remove: Object is not subtype of DeleteOnly
}
```

The translation adds a dummy `remove` public method in `FixedList`. The translated version of the above example is as follows:

```
interface List<E,M> extends Collection<E,M> {
    E remove(int index);
    ...
}
```

```
class FixedList<E> implements List<E,Object> {
    ...
}
```

```
public E remove(int index)
    throws UnsupportedOperationException
{ throw new UnsupportedOperationException(); }
}
```

The same translation technique is used for safety: in a naive translation scheme, subclasses of a class that has conditional methods would inherit those methods because of the erasure translation in the superclass, allowing code not compiled with the cJ compiler to gain access to those methods. Instead, we ensure that the subclass overrides the method with a dummy implementation to avoid such accidental exposure of the superclass functionality. Note that this problem is similar to that faced by the designers of GJ [3], and the solution we adopt is also similar to theirs. For instance, consider the following class:

```
class Channel<T> {
    <T extends Trusted>?
    void disconnect() { ... }
}
```

Erasure will remove the type-conditional and, thus, expose the `disconnect` method. If the user wants to ensure security he/she can export only specialized subclasses that explicitly do not implement the `Trusted` interface:

```
class NonsecureChannel extends Channel<Object> { }
```

The cJ compiler will translate the latter into a class that is safe to use in an insecure environment, avoiding accidental exposure of the superclass method:

```
class NonsecureChannel extends Channel<Object> {
    void disconnect()
    throws UnsupportedOperationException
    { throw new UnsupportedOperationException(); }
}
```

This translation technique does not help avoid the accidental exposure of fields, however. To protect a conditional field against unauthorized access, a programmer could designate the field `private`, and define getter/setter methods for it. The above translation technique for methods can then be used to protect the getter/setter methods from unauthorized uses.

In certain situations, a type instantiation considered legal by the cJ compiler might not be considered legal by a regular Java compiler. For example,

```
class C<X> {
    <X extends Enum<X>>?
    EnumSet<X> es = null;
}
```

```
class EnumSet<X extends Enum<X>> {
    ...
}
```

A simple erasure applied to class `C<X>` would erase the type condition `<X extends Enum<X>>?`. However, type instantiation `EnumSet<X>` is not compilable using a Java compiler, because `X` is nowhere declared to be a subtype of `Enum<X>`. In these situations, cJ translates type `EnumSet<X>` all the way down to its “raw type” form, `EnumSet`. Thus, the translation of `C<X>` would be:

```
class C<X> {
    EnumSet es = null;
}
```

The cJ compiler then needs to perform all the translations that a regular Java compiler does for expressions of type `EnumSet`, e.g., generating casts of return types of methods called on this type.

Translation and Backward Compatibility. The interesting aspect of the cJ translation, as described above, is that it is remarkably

simple and fits very well the existing Java object model. The restrictions of the cJ language outlined in Section 2.2 are in place explicitly so that an elegant erasure-based translation can be supported. For instance, ensuring that methods do not conflict, even when they are under disjoint type conditions, means that we can employ the erasure-based translation without the need for method renamings. Similarly, ensuring that overriding methods (in a subclass) are declared under weaker type conditions than the overridden methods (in the superclass) enables a clean erasure by just removing the type-conditionals. It means that a subclass method does not “accidentally” override a valid superclass method when the subclass method should not really exist based on its type condition. Translating all cJ classes and methods one-to-one into Java classes and methods ensures good interfacing with client code, and even unsuspecting legacy (i.e., standard Java) code.

The cJ translation also includes some transparent special case handling purely for strong backward compatibility, even at the source level. This was motivated by our study of the Java Collections Framework. The special handling occurs when the cJ compiler is invoked in “compatibility mode” and when a type parameter is used *only* in type-conditionals (and not, for instance, to declare references). In that case, the cJ compiler treats the parameter as optional. For instance, the cJ compiler can compile legacy Java code using the `Collection<E>` interface (and any of the classes implementing it) against the cJ library, which defines `Collection` as `Collection<E,M>`. (Either all optional parameters or none need to be omitted.) When a type parameter is omitted, the instantiation is assumed to satisfy all the type-conditionals, and any instantiation with full type parameters is a subtype of it, and vice versa. Our treatment is directly analogous to “raw types” in the translation of Java generics [3].

Furthermore, when a type parameter to a class is used only in type-conditionals (or transitively as a type argument to another class that uses this parameter only in type conditionals), then the cJ compiler removes it from the translated code. This means that the code generated from the cJ compiler can be used as a regular Java library, under plain Java compilers. This is best illustrated with an example. Consider the form of our standard `List` interface from the Java Collections API:

```
interface List<E,M> extends Collection<E,M> {
  <M extends Modifiable?>
  E set(int index, E element);

  <M extends Resizable?>
  void add(int index, E element);
  ...
}
```

`List` uses its type parameter `M` only in type conditions and to instantiate another type, `Collection`, where it is also used only in type conditions. `M` is never used as an argument or return type of a method. Therefore, the cJ compiler accepts code that refers to `List` with only one type parameter. At the same time, the translation of the above cJ interface into a plain Java interface eliminates the second type parameter:

```
interface List<E> extends Collection<E> {
  E set(int index, E element);

  void add(int index, E element);
  ...
}
```

In short, the cJ compiler compiles old-style Java code even against new-style (cJ) libraries that use extra type variables for type conditions. Furthermore, the cJ compiler translates new-style (type conditional) library code into Java code that is source-compatible with existing Java client code, under standard Java compilers.

Clearly, our erasure translation has the same requirements as other erasure translations—e.g., that of GJ [3]—for the purpose of full integration in Java. For instance, the reflection mechanism needs to change to support cJ-translated code. This is not part of our current implementation.

6. Formalization

We present the formal syntax and typing rules for a subset of cJ. Our formalism is an extension of the formalism for Featherweight GJ (FGJ) with variance, by Igarashi and Viroli [12]. We call our calculus Featherweight cJ (FCJ). FCJ captures a core subset of cJ functionality that allows us to explore the type-safety issues introduced by type-conditionals, with minimum extra baggage and duplication of work that has already been done for FGJ [11] and variance [12]. Our formalism requires that all type parameters declare upper bounds, which may be `Object`. Each class must declare a superclass, which may also be `Object`. Additionally, all superclass declarations must be guarded by type-conditionals, though unconditional superclasses can be expressed by having the type-conditional be `<X extends N?>`, where `N` is `X`’s declared upper bound. Similarly, all method declarations must be guarded by type-conditionals, as well. Conditional fields are not supported in the formalism, since the issues involving conditional member declarations are thoroughly represented by conditional methods. Interfaces are not part of either the original FGJ, or our formalism. Thus, we only support conditional superclasses. A class declaration includes a sequence of fields and method declarations. We assume an implicit constructor for each class, which takes as arguments expressions that can be used to initialize field values. The method body is simply an expression.

Note that the variance formalization by Igarashi and Viroli does not strictly model the wildcard implementation in Java. Some notable differences include the inability in the Igarashi and Viroli system to access a co-variantly typed field from a contra-variantly instantiated type, yielding `Object` as the field’s type. An attempt to formalize the wildcard mechanism as it is implemented in Java is presented by Torgersen et al. [28]. However, the Torgersen et al. formalism has not been proven sound. Thus, we choose to work with the Igarashi and Viroli formalization here, as a solid basis for proving the soundness of our type system.⁵

Notation. For readers unfamiliar with FGJ [11] and the variance formalism [12], we briefly introduce the notational conventions used. The meta-variables `C` and `D` range over class names; `X`, `Y`, and `Z` range over distinct type variables; `S`, `T`, `U`, `V`, and `W` range over types; `H`, `N`, `O`, `P`, `Q`, and `R` range over nonvariable types (fully instantiated types); `f` and `g` range over field names; `m` ranges over method names; `x` ranges over parameter names; `d` and `e` range over expressions; and `M` ranges over method declarations. Meta-variable `v` represents variance annotations `o`, `+`, `-`, and `*`, for in-variant, co-variant, contra-variant, and bi-variant, respectively—e.g., `+T` corresponds to `? extends T` in the full Java syntax. Variance annotations can be placed in front of any non-variable type. In-variant is the assumed default, and thus, `C<oT>` is abbreviated to `C<T>`. A partial order \leq on variance annotations can be defined as: $o \leq + \leq *$, $o \leq - \leq *$. $v_1 \vee v_2$ represents the least upper bound of v_1 and v_2 .

In addition, we use a few shorthand conventions for conciseness. \bar{X} is a shorthand for X_1, \dots, X_n , and similarly, $\bar{T} \bar{x}$ is a shorthand for $T_1 x_1, \dots, T_n x_n$. When this shorthand is applied to a type variable or a regular variable (i.e., fields, method arguments), it

⁵ After the completion of the work presented in this paper, an even more recent formalization of Java wildcards has been published [4]. This formalization does have a proof of soundness, and should reflect more accurately the wildcard mechanism in Java. We intend to explore using this formalization in the cJ type system in our future work.

represents a sequence with no duplication. We use \bullet to denote an empty sequence. The notation \triangleleft is the shorthand for keyword `extends`, and \uparrow is the shorthand for keyword `return` in method bodies.

We also assume a class table CT , which maps class names C to their declarations. A *program* is a pair (CT, e) of a class table, and an expression.

6.1 Syntax

We present the FCJ syntax in Figure 1. The syntax follows closely the abstract syntax for FGJ with variance [12]. The main difference is the addition of a type-conditional construct in front of superclass and method declarations. The type-conditional construct, $\langle \bar{X} \triangleleft \bar{R} \rangle ?$, evaluates to true if, after type parameter instantiation, the types for \bar{X} are subtypes of \bar{R} . A fully instantiated class has the declared superclass if and only if the type-conditional guarding the superclass declaration evaluates to true. Otherwise, it extends `Object`. Similarly, a method exists for a fully instantiated class if and only if its type condition evaluates to true.

Syntax:	
T	$::= X \mid N$
N	$::= C \langle \bar{V} \bar{T} \rangle$
v	$::= o \mid + \mid - \mid *$
CL	$::= \text{class } C \langle \bar{X} \triangleleft \bar{N} \rangle \text{ S-if } \langle D \triangleleft \bar{S} \rangle \{ \bar{T} \bar{f}; \bar{M} \}$
M	$::= \text{S-if } \langle \bar{Y} \triangleleft \bar{P} \rangle T \text{ m } (\bar{T} \bar{x}) \{ \uparrow e; \}$
S-if	$::= \langle \bar{X} \triangleleft \bar{R} \rangle ?$
e	$::= x$
	$\mid e.f$
	$\mid e. \langle \bar{T} \rangle m(\bar{e})$
	$\mid \text{new } C \langle \bar{T} \rangle (\bar{e})$
	$\mid (T)e$

Figure 1. Syntax

6.2 Type System

The main typing rules for FCJ are presented in Figure 2. Δ and Γ are the two environments used in typing judgments. Δ is a type environment that ranges over subtyping assumptions of the form $T \triangleleft S$. When $X \triangleleft N \in \Delta$ and N is a non-variable type, for all X , we say that Δ has non-variable bounds. Γ is a variable environment that maps a variable x to its type T .

To support the typing rules, we present some auxiliary definitions in Figure 3, and the definition of “open” (\uparrow^Δ) and “close” (\downarrow_Δ) of variant types in Figure 4. These rules and definitions follow closely the format of those in variance-based FGJ. We assume the reader has a certain familiarity with the FGJ formalization, though not necessarily with the Igarashi and Viroli variance formalism. To enhance the understanding of our type system, we first highlight some important additions to FGJ made by Igarashi and Viroli regarding variance. We then delve into the rules and definitions specifically changed for the inclusion of type-conditionals in cJ.

Background on Variance Formalism. The two most important additions of the Igarashi and Viroli system over FGJ are the concepts of “open” and “close” (Figure 4). Before any type is used (i.e., for field or method invocation, or in a subtyping judgment), it must be “opened” first. Opening a type means that we introduce a fresh type variable for each co- or contra-variantly defined type. For example, before we can check the validity of invoking method m in type $C \langle +T \rangle$, we must open this type by introducing a fresh type variable X into Δ , where $\Delta \vdash X \triangleleft T$. To look for method m in

$C \langle +T \rangle$ now means to look for m in $C \langle X \rangle$. If T occurs anywhere in m ’s type, it is replaced by X , as well.

This “opening” conveniently disallows illegal accesses of methods or fields that we informally described in Section 4. For example, suppose that in the definition of class $C \langle X \rangle$, we have method $D \text{ m } (X \ x) \{ \dots \}$. The type parameter X appears in a contra-variant position—as method m ’s argument type. This means that any co-variantly instantiated type $C \langle +T \rangle$ should not be able to invoke method m . This is indeed the case in this formalism: we first open $C \langle +T \rangle$ to $C \langle Y \rangle$, where $\Delta \vdash Y \triangleleft T$. We then check that any invocation of m passes in an argument of some subtype of Y . However, Y is simply a type variable with an *upper* bound of T . According to the subtyping rules in Figure 2, no type can be deemed a subtype of Y (except Y itself, which is not available before the opening, and thus cannot be the type of any argument passed to m). Thus, no invocation of m on an expression of type $C \langle +T \rangle$ can be well-typed. Similarly, had the type parameter X appeared in a co-variant position in $C \langle X \rangle$, e.g., as the return type of a method, an expression with the contra-variantly instantiated type $C \langle -T \rangle$ would not have been able to invoke that method.

Since “open” introduces new type variables into the type environment, it is always paired with a “close” operation, where the newly introduced type variable is closed down to its bound and removed from the type environment. Closing also re-introduces variance annotations, using a conservative combination of the variance annotations of the current use context (i.e., the type being closed) and the surrounding definition context used for the preceding “open”.

Auxiliary Definitions. Function $mtype(\Delta, m, C \langle \bar{T} \rangle)$ returns the signature of method m , in type $C \langle \bar{T} \rangle$, in the form of $\langle \bar{Y} \triangleleft \bar{P} \rangle \bar{U} \rightarrow U_0$. $mtype(\Delta, m, C \langle \bar{T} \rangle)$ is defined under two rules:

- **MT-CLASS** says that if method m is declared in class $C \langle \bar{X} \rangle$ with type-conditional $\langle \bar{X} \triangleleft \bar{R} \rangle ?$, and the type-conditional is satisfied by substituting types \bar{T} for type parameters \bar{X} , i.e., $\Delta \vdash \bar{T} \triangleleft \bar{R}$, then $mtype(\Delta, m, C \langle \bar{T} \rangle)$ is defined.
- **MT-SUPER** covers the condition when method m is not declared in class $C \langle \bar{X} \rangle$ at all. In this case, if the type-conditional for superclass $D \langle \bar{S} \rangle$ is satisfied by the substitution $\frac{\bar{T}}{\bar{X}}$, then $mtype(\Delta, m, C \langle \bar{T} \rangle)$ is defined as $mtype(\Delta, m, \frac{\bar{T}}{\bar{X}}[D \langle \bar{S} \rangle])$.

Note that we do not need a case for when m is declared in $C \langle \bar{X} \rangle$, but the type-conditional guarding it is not satisfied by the substitution $\frac{\bar{T}}{\bar{X}}$. As explained in Section 2.2, the type conditions on a subclass method must be weaker than the type conditions guarding the method it overrides in the superclass (this restriction is formalized in the *override* rule, which we explain later in this section). Thus, if method m ’s type conditions in $C \langle \bar{X} \rangle$ cannot be satisfied by the assumptions in Δ , then its type conditions in the superclass of $C \langle \bar{X} \rangle$ cannot possibly be satisfied. There is no need to invoke MT-SUPER in this case.

$mbody(\Delta, m \langle \bar{w} \rangle, C \langle \bar{T} \rangle)$ returns a pair, (\bar{x}, e) . \bar{x} are the parameters of m , and e is m ’s body. \bar{w} are the actual types inferred for a polymorphic method m . Note $mbody$ is similarly defined under the same two conditions that $mtype$ is.

$fields(\Delta, C \langle \bar{T} \rangle)$ returns a sequence of fields in class $C \langle \bar{T} \rangle$. `Object` has no fields. For all other types $C \langle \bar{T} \rangle$, $fields(\Delta, C \langle \bar{T} \rangle)$ returns the sequence of fields declared in $C \langle \bar{T} \rangle$, and, if the type-conditional guarding $C \langle \bar{T} \rangle$ ’s superclass, $D \langle \bar{S} \rangle$, is satisfied by the substitution $\frac{\bar{T}}{\bar{X}}$, the value of $fields(\Delta, \frac{\bar{T}}{\bar{X}}[D \langle \bar{S} \rangle])$ is returned, as well.

The predicate $override(\Delta, m, \langle \bar{X} \triangleleft \bar{R} \rangle ? D \langle \bar{S} \rangle, \langle \bar{X} \triangleleft \bar{H} \rangle ? \langle \bar{Y} \triangleleft \bar{P} \rangle \bar{U} \rightarrow U_0 \rightarrow U_0)$ judges if a method m , with signature $\langle \bar{X} \triangleleft \bar{H} \rangle ? \langle \bar{Y} \triangleleft \bar{P} \rangle \bar{U} \rightarrow U_0$ may be defined in a class that has a conditional superclass $D \langle \bar{S} \rangle$,

Expression typing:	Class typing:	
$\frac{\Delta; \Gamma \vdash x \in \Gamma(x)}{\quad} \quad (T\text{-VAR})$	$\frac{\Delta_1 = \bar{X} <: \bar{N} \quad \Delta_1 \vdash \bar{R} <: \bar{N} \quad \Delta_1 \vdash \bar{N} \text{ ok} \quad \Delta_2 = \bar{X} <: \bar{R} \quad \Delta_2 \vdash \bar{R}, D < \bar{S} > \text{ ok} \quad \Delta_1 \vdash \bar{T} \text{ ok} \quad \bar{M} \text{ OK IN } C < \bar{X} < \bar{N} >}{\text{class } C < \bar{X} < \bar{N} > < \bar{X} < \bar{R} > ? < D < \bar{S} > \{ \bar{T} \bar{f}; \bar{M} \} \text{ OK}} \quad (T\text{-CLASS})$	
$\frac{\Delta; \Gamma \vdash e_0 \in T_0 \quad \Delta \vdash \text{bound}_{\Delta}(T_0) \uparrow \Delta' C < \bar{U} > \quad \text{fields}(\Delta, \text{bound}_{\Delta}(C < \bar{U} >)) = \bar{S} \bar{f} \quad S_i \downarrow_{\Delta'} T}{\Delta; \Gamma \vdash e_0.f_i \in T} \quad (T\text{-FIELD})$	Subtyping:	
$\frac{\Delta; \Gamma \vdash e_0 \in T_0 \quad \Delta \vdash \text{bound}_{\Delta}(T_0) \uparrow \Delta' C < \bar{T} > \quad \text{mtype}(\Delta, m, C < \bar{T} >) = < \bar{Y} < \bar{P} > \bar{U} \rightarrow U_0 \quad \bar{Y} <: \bar{Q} \notin \Delta' \text{ for any } \bar{Q} \quad \Delta \vdash \bar{V} \text{ ok} \quad \Delta, \Delta' \vdash \bar{V} <: \bar{Y} < \bar{Y} > \bar{P} \quad \Delta; \Gamma \vdash \bar{\theta} \in \bar{S} \quad \Delta, \Delta' \vdash \bar{S} <: \bar{Y} < \bar{Y} > \bar{U} \quad \bar{Y} < \bar{Y} > U_0 \downarrow_{\Delta'} T}{\Delta; \Gamma \vdash e_0.< \bar{V} >_m(\bar{\theta}) \in T} \quad (T\text{-INVK})$	$\frac{\Delta \vdash T <: T}{\Delta \vdash T <: T} \quad (S\text{-REFL})$	
$\frac{\Delta \vdash C < \bar{T} > \text{ ok} \quad \text{fields}(\Delta, C < \bar{T} >) = \bar{U} \bar{f} \quad \Delta; \Gamma \vdash \bar{\theta} \in \bar{S} \quad \Delta \vdash \bar{S} <: \bar{U}}{\Delta; \Gamma \vdash \text{new } C < \bar{T} >(\bar{\theta}) \in C < \bar{T} >} \quad (T\text{-NEW})$	$\frac{\Delta \vdash S <: T \quad \Delta \vdash T <: U}{\Delta \vdash S <: U} \quad (S\text{-TRANS})$	
$\frac{\Delta; \Gamma \vdash e_0 \in T_0 \quad \Delta \vdash T \text{ ok} \quad \Delta \vdash \text{bound}_{\Delta}(T_0) <: \text{bound}_{\Delta}(T) \quad \text{or } \Delta \vdash \text{bound}_{\Delta}(T) <: \text{bound}_{\Delta}(T_0)}{\Delta; \Gamma \vdash (T)e_0 \in T} \quad (T\text{-CAST})$	$\frac{\frac{X <: T \in \Delta}{\Delta \vdash X <: T} \quad \frac{T <: X \in \Delta}{\Delta \vdash T <: X}}{CT(C) = \text{class } C < \bar{X} < \bar{N} > < \bar{X} < \bar{R} > ? < D < \bar{S} > \{ \dots \} \quad \Delta \vdash C < \bar{T} > \uparrow \Delta' C < \bar{U} > \quad \Delta, \Delta' \vdash \bar{U} <: \bar{U} < \bar{X} > \bar{R} \quad ((\bar{U} < \bar{X} > D < \bar{S} >) \downarrow_{\Delta'} T)}{\Delta \vdash C < \bar{T} > <: T} \quad (S\text{-UBOUND})$	
$\frac{\Delta; \Gamma \vdash e_0 \in T_0 \quad \Delta \vdash T \text{ ok} \quad \Delta \not\vdash \text{bound}_{\Delta}(T_0) <: \text{bound}_{\Delta}(T) \quad \text{and } \Delta \not\vdash \text{bound}_{\Delta}(T) <: \text{bound}_{\Delta}(T_0)}{\Delta; \Gamma \vdash (T)e_0 \in T} \quad (T\text{-SCAST})$	$\frac{\bar{v} \leq \bar{w} \quad \text{if } w_i \leq -, \text{ then } \Delta \vdash T_i <: S_i \quad \text{if } w_i \leq +, \text{ then } \Delta \vdash S_i <: T_i}{\Delta \vdash C < \bar{v} > <: C < \bar{w} >} \quad (S\text{-LBOUND})$	
Method typing:	$\frac{\bar{X} <: \bar{N} < \bar{R} <: \bar{N} \quad \Delta = \bar{X} <: \bar{R}, \bar{Y} <: \bar{P} \quad \Delta \vdash \bar{R}, \bar{P}, \bar{T}, T_0 \text{ ok} \quad \Delta; \bar{x} : \bar{T}, \text{this} : C < \bar{X} > \vdash e_0 \in S_0 \quad \Delta \vdash S_0 <: T_0 \quad CT(C) = \text{class } C < \bar{X} < \bar{N} > < \bar{X} < \bar{R} > ? < D < \bar{S} > \{ \dots \} \quad \text{override}(\Delta, m, < \bar{X} < \bar{V} > ? D < \bar{S} >, < \bar{X} < \bar{R} > ? < \bar{Y} < \bar{P} > \bar{T} \rightarrow T_0)}{\frac{< \bar{X} < \bar{R} > ? < \bar{Y} < \bar{P} > T_0 \text{ m } (\bar{T} \bar{x}) \{ \uparrow e_0; \} \text{ OK IN } C < \bar{X} < \bar{N} >}{\quad}} \quad (T\text{-METHOD})$	Well-formed types:
	$\frac{\Delta \vdash \text{Object ok}}{\Delta \vdash \text{Object ok}} \quad (WF\text{-OBJECT})$	
	$\frac{X <: T \in \Delta}{\Delta \vdash X \text{ ok}} \quad (WF\text{-VAR})$	
	$\frac{CT(C) = \text{class } C < \bar{X} < \bar{N} > < \bar{X} < \bar{R} > ? < D < \bar{S} > \{ \dots \} \quad \Delta \vdash \bar{T} \text{ ok} \quad \Delta \vdash \bar{T} <: \bar{T} < \bar{X} > \bar{N}}{\Delta \vdash C < \bar{T} > \text{ ok}} \quad (WF\text{-CLASS})$	

Figure 2. Typing Rules

guarded by condition $\langle \bar{X} < \bar{R} \rangle$. The extra complication in this rule reflects the restrictions described in Section 2.2, and is used in proving the correctness of our erasure-based translation (Section 6.4). There are two aspects of our translation to consider. Firstly, recall that our translation scheme erases all type conditionals. After translation, a class $C < \bar{X} >$ extends its superclass $D < \bar{S} >$ unconditionally. Consequently, even if a method m in $C < \bar{X} >$ is declared under a type-conditional that precludes the condition for the superclass, the type of m in $C < \bar{X} >$ still cannot conflict with the type of m in $D < \bar{S} >$. Secondly, also recall that if a method in a subclass has the same type signature as a method in a superclass, we require the subclass method to *always* override the superclass method. This means the type-conditional on the subclass method must be implied by the type-conditional on the superclass method. This requirement ensures that we do not have to dynamically decide whether a class's own method implementation should be invoked or it should call `super.m(...)`.

In order for *override* to reflect these restrictions, it uses the function mtype_{uc} , which *unconditionally* recurses up the chain of superclasses to find a method's signature. $\text{mtype}_{uc}(m, C < \bar{T} >)$ returns a pair, $(\Delta', < \bar{Y} < \bar{P} > \bar{U} \rightarrow U_0)$. Δ' contains subtyping assumptions that must be satisfied for method m to have type $< \bar{Y} < \bar{P} > \bar{U} \rightarrow U_0$. The second part of the pair is m 's signature as defined in the closest superclass up the unconditional chain of inheritance.

The *override* rule uses the *implies* notation (as in FGJ [11]) to indicate that the restrictions represented by the consequent of the *implies* need to be satisfied only if the antecedent is true. In

this case, the antecedent is: $\text{mtype}_{uc}(m, \bar{X}, D < \bar{S} >) = (\Delta', < \bar{Z} < \bar{Q} > \bar{T} \rightarrow T_0)$. This means that method m is defined in either $D < \bar{S} >$, or some conditional superclass of $D < \bar{S} >$. If this antecedent is true, then the parameter, return types, and bounds on the inferred types must be the same in the subclass as they are in the conditional superclass. It must also be true that given all conditions guarding the chain of conditional superclasses mtype_{uc} recursed through to find m , and the condition guarding m itself, the condition guarding the subclass method is true, as well. This is checked by augmenting Δ with Δ' and $\bar{X} <: \bar{R}$, and requiring that these are sufficient to show $\bar{X} <: \bar{R}$.

Note that the definition of *override* is dependent on the subtyping rules defined in Figure 2. Depending on the specific algorithm implementing our declarative subtyping rules, it is possible that the subtyping condition guarding the overriding method, $\bar{X} <: \bar{R}$, cannot be shown to be true using the assumptions in Δ , Δ' , and $\bar{X} <: \bar{R}$. Consequently, certain valid overriding methods cannot be proven so. To see this concretely, let $\Delta = \text{Foo} < X > <: \text{Baz}$, where `Foo` is defined as: `class Foo < X > < X < Bar > ? < Baz > { ... }`. If we want to show that $\Delta \vdash X <: \text{Bar}$, we need to deconstruct type `Foo < X >`, and infer from `Foo < X > <: Baz` that $X <: \text{Bar}$ must be true, as well. Our current implementation does not deconstruct types to do such inference. Subtyping assumptions thrown into Δ' by mtype_{uc} are only effective if the type variables \bar{X} are not buried inside of constructed types, such as `Foo < X >`. We are currently working on a decidable algorithm for deconstructing types to get more precise subtyping assumptions. Note that this is a standard point of trade-off: a too powerful reasoning procedure may well end up being undecidable.

<p>Method type lookup:</p> $\frac{CT(C) = \text{class } C \langle \bar{X} \langle \bar{N} \rangle \langle \bar{X} \langle \bar{Q} \rangle ? \langle D \langle \bar{S} \rangle \{ \dots \bar{M} \} \rangle \rangle \langle \bar{Y} \langle \bar{P} \rangle U_0 \ m \ (\bar{U} \ \bar{x}) \ \{ \uparrow e; \} \in \bar{M} \quad \Delta \vdash \bar{T} \langle : [\bar{T} / \bar{X}] \bar{R} \rangle}{mtype(\Delta, m, C \langle \bar{T} \rangle) = [\bar{T} / \bar{X}] (\langle \bar{Y} \langle \bar{P} \rangle \bar{U} \rightarrow U_0)} \quad (\text{MT-CLASS})}$ $\frac{CT(C) = \text{class } C \langle \bar{X} \langle \bar{N} \rangle \langle \bar{X} \langle \bar{Q} \rangle ? \langle D \langle \bar{S} \rangle \{ \dots \bar{M} \} \rangle \rangle \ m \ \text{is not defined in } \bar{M} \quad \Delta \vdash \bar{T} \langle : [\bar{T} / \bar{X}] \bar{Q} \rangle}{mtype(\Delta, m, C \langle \bar{T} \rangle) = mtype(\Delta, m, [\bar{T} / \bar{Q}] D \langle \bar{S} \rangle)} \quad (\text{MT-SUPER})}$ <p>Method body lookup:</p> $\frac{CT(C) = \text{class } C \langle \bar{X} \langle \bar{N} \rangle \langle \bar{X} \langle \bar{Q} \rangle ? \langle D \langle \bar{S} \rangle \{ \dots \bar{M} \} \rangle \rangle \langle \bar{X} \langle \bar{R} \rangle ? \langle \bar{Y} \langle \bar{P} \rangle U_0 \ m \ (\bar{U} \ \bar{x}) \ \{ \uparrow e; \} \in \bar{M} \quad \Delta \vdash \bar{T} \langle : [\bar{T} / \bar{X}] \bar{R} \rangle}{mbody(\Delta, m \langle \bar{W} \rangle, C \langle \bar{T} \rangle) = (\bar{x}, [\bar{W} / \bar{Y}] [\bar{T} / \bar{X}] e)} \quad (\text{MB-CLASS})}$ $\frac{CT(C) = \text{class } C \langle \bar{X} \langle \bar{N} \rangle \langle \bar{X} \langle \bar{Q} \rangle ? \langle D \langle \bar{S} \rangle \{ \bar{S} \ \bar{f}; \bar{M} \} \rangle \rangle \ m \ \text{is not defined in } \bar{M} \quad \Delta \vdash \bar{T} \langle : [\bar{T} / \bar{X}] \bar{Q} \rangle}{mbody(\Delta, m \langle \bar{W} \rangle, C \langle \bar{T} \rangle) = mbody(\Delta, m \langle \bar{W} \rangle, [\bar{T} / \bar{X}] D \langle \bar{S} \rangle)} \quad (\text{MB-SUPER})}$ <p>Field lookup:</p> $fields(\Delta, \text{Object}) = \bullet$ $\frac{CT(C) = \text{class } C \langle \bar{X} \langle \bar{N} \rangle \langle \bar{X} \langle \bar{R} \rangle ? \langle D \langle \bar{U} \rangle \{ \bar{S} \ \bar{f}; \bar{M} \} \rangle \rangle \quad fields(\Delta, [\bar{T} / \bar{X}] D \langle \bar{U} \rangle) = \bar{D} \ \bar{g} \quad \Delta \vdash \bar{T} \langle : [\bar{T} / \bar{X}] \bar{R} \rangle}{fields(\Delta, C \langle \bar{T} \rangle) = \bar{D} \ \bar{g}, [\bar{T} / \bar{X}] \bar{S} \ \bar{f}}$ $\frac{CT(C) = \text{class } C \langle \bar{X} \langle \bar{N} \rangle \langle \bar{X} \langle \bar{R} \rangle ? \langle D \langle \bar{U} \rangle \{ \bar{S} \ \bar{f}; \bar{M} \} \rangle \rangle \quad \Delta \vdash \bar{T} \langle : [\bar{T} / \bar{X}] \bar{R} \rangle}{fields(\Delta, C \langle \bar{T} \rangle) = [\bar{T} / \bar{X}] \bar{S} \ \bar{f}}$	<p>Unconditional field lookup:</p> $fields_{uc}(\text{Object}) = \bullet$ $\frac{CT(C) = \text{class } C \langle \bar{X} \langle \bar{N} \rangle \langle \bar{X} \langle \bar{R} \rangle ? \langle D \langle \bar{U} \rangle \{ \bar{S} \ \bar{f}; \bar{M} \} \rangle \rangle \quad fields_{uc}([\bar{T} / \bar{X}] D \langle \bar{U} \rangle) = \bar{D} \ \bar{g}}{fields_{uc}(C \langle \bar{T} \rangle) = \bar{D} \ \bar{g}, [\bar{T} / \bar{X}] \bar{S} \ \bar{f}}$ <p>Unconditional method type lookup:</p> $\frac{CT(C) = \text{class } C \langle \bar{X} \langle \bar{N} \rangle \langle \bar{X} \langle \bar{Q} \rangle ? \langle D \langle \bar{S} \rangle \{ \dots \bar{M} \} \rangle \rangle \langle \bar{X} \langle \bar{R} \rangle ? \langle \bar{Y} \langle \bar{P} \rangle U_0 \ m \ (\bar{U} \ \bar{x}) \ \{ \uparrow e; \} \in \bar{M} \quad \Delta' = [\bar{T} / \bar{X}] (\bar{X} \langle : \bar{R} \rangle)}{mtype_{uc}(m, C \langle \bar{T} \rangle) = (\Delta', [\bar{T} / \bar{X}] (\langle \bar{Y} \langle \bar{P} \rangle \bar{U} \rightarrow U_0))} \quad (\text{MT}_{UC}\text{-CLASS})}$ $\frac{CT(C) = \text{class } C \langle \bar{X} \langle \bar{N} \rangle \langle \bar{X} \langle \bar{Q} \rangle ? \langle D \langle \bar{S} \rangle \{ \dots \bar{M} \} \rangle \rangle \quad m \ \text{is not defined in } \bar{M}}{mtype_{uc}(m, D \langle \bar{S} \rangle) = (\Delta, \langle \bar{Y} \langle \bar{P} \rangle \bar{U} \rightarrow U_0) \quad \Delta' = \Delta, [\bar{T} / \bar{X}] (\bar{X} \langle : \bar{Q} \rangle)} \quad (\text{MT}_{UC}\text{-SUPER})}$ <p>Valid method overriding:</p> $\frac{mtype_{uc}(m, D \langle \bar{S} \rangle) = (\Delta', \langle \bar{Z} \langle \bar{Q} \rangle \bar{T} \rightarrow T_0) \ \text{implies} \quad [\bar{Y} / \bar{Z}] (\bar{T}, T_0, \bar{Q}) = (\bar{U}, U_0, \bar{P}) \ \text{and} \ \Delta, \Delta', \bar{X} \langle : \bar{R} \rangle \vdash \bar{X} \langle : \bar{H} \rangle}{override(\Delta, m, \langle \bar{X} \langle \bar{R} \rangle ? D \langle \bar{S} \rangle, \langle \bar{X} \langle \bar{H} \rangle ? \langle \bar{Y} \langle \bar{P} \rangle \bar{U} \rightarrow U_0 \rangle)}$ <p>Bound of type:</p> $bound_{\Delta}(N) = N \quad \frac{\Delta(X) = (+, S)}{bound_{\Delta}(X) = bound_{\Delta}(S)}$
--	--

Figure 3. Auxiliary definitions

A conservative algorithm, on the other hand, will reject some programs because of its inability to establish the conditions for their soundness. The latter is typically preferable in practice, since, in this setting, troublesome programs tend to be highly contrived.

The two rules for unconditional field lookup are only used in proving the correctness of erasure, using the erasure rules detailed in the accompanying technical report [10]. They are included for completeness of the auxiliary functions.

Type Rules. Most of the rules presented in Figure 2 are the same as their variance-based FGJ counterparts. We now go through the ones particular to the type-conditional extensions in FCJ.

T-FIELD and **T-INVK**: these rules define when a field reference or a method invocation, respectively, is well-typed. Even though they look identical to their variance-based FGJ counterparts, they use functions *fields* and *mtype*, which fully encapsulate the lookup of conditional supertypes and conditional methods, as previously explained.

T-METHOD: The interesting change to the T-METHOD rule from its counterpart in variance-based FGJ is that the environment Δ (under which the return and parameter types, as well as the body of the method, expression e_0 , must be well-typed) is augmented with the type bound $\bar{X} \langle : \bar{R} \rangle$, which is the type-conditional under which the method is declared. Intuitively, this says that if a method is declared under type-conditional $\langle \bar{X} \langle \bar{R} \rangle ?$, then in the scope of the body of the method it can be assumed that the type environment Δ supports this bound.

T-CLASS: Note that the conditional superclass $D \langle \bar{S} \rangle$ needs to be proved well-typed under Δ augmented with the type-conditional condition guarding it.

Open:	$\Delta \vdash T \uparrow^{\emptyset} T \quad (\text{O-REFL})$ $\frac{\Delta \vdash S \uparrow^{\Delta_1} T \quad \Delta, \Delta_1 \vdash T \uparrow^{\Delta_2} U}{\Delta \vdash S \uparrow^{\Delta_1, \Delta_2} U} \quad (\text{O-TRANS})$ $\frac{X \ \text{fresh for } \Delta, C \langle \bar{v}_1 \bar{T}_1, vT, \bar{v}_2 \bar{T}_2 \rangle \quad v \neq o}{\Delta \vdash C \langle \bar{v}_1 \bar{T}_1, vT, \bar{v}_2 \bar{T}_2 \rangle \uparrow^{X: (v, T)} C \langle \bar{v}_1 \bar{T}_1, oX, \bar{v}_2 \bar{T}_2 \rangle} \quad (\text{O-CLASS})$
Close:	$\frac{\Delta(X) = (+, T)}{X \downarrow_{\Delta} T} \quad (\text{C-PROM})$ $\frac{X \notin dom(\Delta)}{X \downarrow_{\Delta} X} \quad (\text{C-TVAR})$ $(w_i, T'_i) = \begin{cases} (v_i, T_i) & \text{if } T_i \downarrow_{\Delta} T_i \\ (v_i \vee +, U_i) & \text{if } T_i \downarrow_{\Delta} U_i \ \text{and } T_i \neq U_i \\ (v_i \vee v'_i, U_i) & \text{if } T_i = X \ \text{and } \Delta(X) = (v'_i, U_i) \end{cases}$ $C \langle \bar{v} \bar{T} \rangle \downarrow_{\Delta} C \langle \bar{w} \bar{T}' \rangle \quad (\text{C-CLASS})$

Figure 4. Open and Close

6.3 Proof of Soundness

We prove the soundness of our type system by proving subject reduction and progress properties [30]. Due to space limitations, we state the theorems here. Interested readers can find the definition of reduction rules, as well as the full version of the proofs in the technical report available on the cJ website [10].

Theorem 1 [Subject Reduction]: If $\Delta; \Gamma \vdash e \in T$ and $e \rightarrow e'$, then $\Delta; \Gamma \vdash e' \in S$ and $\Delta \vdash S \langle : T$ for some S .

Theorem 2 [Progress]: Let e be a well-typed expression.

1. If e has $\text{new } C\langle\bar{T}\rangle(\bar{e}) . f$ as a subexpression, then $\text{fields}(\emptyset, C\langle\bar{T}\rangle) = \bar{U} \bar{f}$, and $\mathbf{f} = \mathbf{f}_i$.

2. If e has $\text{new } C\langle\bar{T}\rangle(\bar{e}) . m(\bar{d})$ as a subexpression, then $\text{mbody}(\emptyset, m, C\langle\bar{T}\rangle) = (\bar{x}, e_0)$ and $|\bar{x}| = |\bar{d}|$.

Theorem 3 [Type Soundness]: If $\emptyset; \emptyset \vdash e \in T$ and $e \rightarrow^* e'$ being a normal form, then e' is either a value v such that $\emptyset; \emptyset \vdash v \in S$ and $\emptyset \vdash S <: T$ for some S , or an expression that includes $(T)\text{new } C\langle\bar{T}\rangle(\bar{e})$ where $\emptyset \vdash C\langle\bar{T}\rangle \not\prec T$.

6.4 Proof of Correctness of Erasure

We formalized our erasure implementation by defining an erasure function that transforms FCJ expressions and types into variance-based FGJ (FGJ_v) expressions and types. $|\cdot|_\Delta$ yields a FGJ_v type by erasing a FCJ type T , and $|e|_{\Delta, \Gamma}$ yields an FGJ_v expression by erasing an FCJ expression e . We prove the following two theorems⁶:

Theorem 4 [Erasure Preserves Typing]: For a program (CT, e) , if CT is ok, and $\Delta; \Gamma \vdash_{FCJ} e \in T$, then $|CT|_\Delta$ is ok, and $|\Delta|_\Delta, |\Gamma|_\Delta \vdash_{FGJ} |e|_{\Delta, \Gamma} \in |T|_\Delta$

Theorem 5 [Erasure Preserves Execution Modulo Expansion]: If $\Delta; \Gamma \vdash_{FCJ} e \in T$ and $e \rightarrow^*_{FCJ} e'$, then there exists some FGJ_v expression d' such that $|e'|_{\Delta, \Gamma} \Rightarrow d'$ and $|e|_{\Delta, \Gamma} \rightarrow^*_{FGJ_v} d'$, where \Rightarrow is the expansion function.

Again, we refer readers to the technical report for definition of the erasure function, expansion function, and proofs of the above two theorems.

7. Related Work

cJ is related to several programming language and software engineering concepts. These range from mainstream modularization techniques to meta-programming and conditional compilation approaches.

Clearly the idea of a type-conditional is closely related to conditional compilation, as with the C/C++ preprocessor “`#ifdef`” construct. Although `#ifdef` is valuable for configuring large projects, it addresses very different needs from cJ. Conditional compilation gives low-level manual control for software configuration. In the context of a portable language, like Java, an `#ifdef` statement becomes less useful. At the same time, conditional compilation suffers from the lack of any form of safety control. The use of conditional flags may be inconsistent, resulting in invalid configurations that are not detected until one attempts to select them. There has been work on adding some safety to conditional compilation by analyzing all configurations of a C program, and there is evidence that such a heuristic approach may work in several contexts—especially for refactoring [8]. Nevertheless, cJ offers full static safety guarantees, eliminating the problem altogether. Furthermore, the type-conditions of cJ are structured, richer than mere propositions, and well-integrated with the Java type system.

Conditional methods have been explored in OO language in work at least as early as CLU [19]. Nevertheless, CLU does not support subtyping, so the language context of this work is notably different. It is, thus, difficult to compare CLU to cJ, where our main goal is to maintain static type safety, yet, at the same time, maintain a clean subtyping hierarchy used for abstraction. Past work on optional methods in Java was also presented by Myers et al. [22]. This was in the context of a proposal for adding genericity to Java, and it includes the feature of attaching `where` clauses to

individual methods. The conditions on the `where` clauses, however, can only be “structural” constraints—i.e., does type parameter T provide method `void foo()`; This mechanism does not support conditional subtyping—e.g., it is not possible to express that a `Collection` is `Comparable`, if the elements it holds are `Comparable`. Even more importantly, the work by Myers et al. does not support type-safe abstraction over classes with conditional methods, as in the interaction of cJ with variance.

More recently, Emir et al. presented an extension to C# to support generalized type constraints on methods [7]. This extension allows both upper and lower bound type conditions on methods. This is similar to cJ in that methods exist conditionally based on the instantiation of parametric types. Nevertheless, there are significant differences, and in future work we plan to pursue combining the two approaches. cJ currently does not allow using the type parameter of a polymorphic method inside a type conditional—a crucial feature in Emir et al.’s work. At the same time, cJ has several features not found in the generalized type constraints approach. First, cJ supports conditional definitions of fields, as well as conditional subtyping. Furthermore, the cJ (and Java) form of variance we examined earlier is a “use-site variance” mechanism as opposed to the “definition-site variance” supported in Emir et al.’s work. In addition to being part of standard Java, we believe that use-site variance is a mechanism better suited for imperative programming languages in general. In this setting, a single class definition is unlikely to produce types that are purely co-variant, purely contra-variant, or purely bi-variant. Instead defining a class will likely implicitly yield a co-variant part, a contra-variant part, etc. In use-site variance these subsets of the class functionality are derived automatically from a single definition. In definition-site variance, they have to be explicitly separated out into distinct interfaces by the programmer. Thus, we believe use-site variance to be a more user-friendly system and a natural fit for Java.

In languages with (multi-)methods outside classes, the work on constraint-based polymorphism in Cecil [20] is related to cJ. Cecil provides users the ability to add constraints to both methods and supertypes. The constraints can be subtyping constraints, as well as structural constraints, requiring a type to provide a particular method. This is a very different context from that of our work, however. Furthermore, the Cecil type system does not have an analogue of our variance approach to abstracting over all objects with or without some of the conditionally defined members.

Our type-conditional is also related to traditional meta-programming techniques, which offer mechanisms for programs to generate other programs. Recent approaches, such as SafeGen [9] and Genoupe [5] attempt to add safety guarantees to meta-programming, yet maintain expressiveness. Nevertheless, these approaches either fail to achieve full safety, or reject programs in a way that is not transparent to the programmer. Neither mechanism integrates seamlessly with a programming language, as cJ does. For instance, SafeGen uses an automatic theorem prover in order to prove well-formedness properties of every produced program. Yet this approach is not guaranteed to always produce accurate results, as the theorem prover may not terminate. Similarly, Genoupe suffers from potential unsafeties, as its reasoning on the safety of generated code relies on the equivalence of arbitrarily complex expressions from the generator source code, which is undecidable to determine. (Based on the published description, it seems that Genoupe unsoundly estimates the run-time equivalence of expressions based on syntactic similarity.)

It is tempting to find parallels between cJ and advanced OO modularization mechanisms such as traits [6, 25], mixins [2], or mixin layers [27]. These approaches vary in expressiveness and several of them are insufficient for solving the combinatorial explosion problems identified in Section 3. For instance, C++-based

⁶For a formal definition of erasure that matches very closely our implementation, the type rules presented in Figure 2 require minor refinement for the proofs of these theorems. These differences are inconsequential to the material in this paper, and are presented in our technical report.

mixins or mixin layers would still require a large number of compositions, with explicit subtyping links added among them, in order to express the required functionality of the Java Collections Framework. It is possible that a traits-based mechanism could serve to alleviate many of the problems in the Java Collections Framework (albeit with a complete rewrite). Nevertheless, there is no such mechanism currently for Java that would tie well with the rest of the language's type system (e.g., variance) and execution model. Furthermore, no mixin or traits mechanism offers capabilities similar to those shown in Section 3.1, i.e., the ability to add extra members only when a type parameter that is already used for other purposes has a certain subtyping property.

Type-conditionals in cJ can be viewed as being similar to type-safe variant records work—e.g., [24]. Nevertheless, variant records mechanisms typically try to address the problem of *run-time* variability with static type-safety. cJ is not concerned with changes to the type of a variable during run-time. Instead, cJ focuses on the static configurability of components. The techniques used to ensure static type safety in the case of variant records and in the case of cJ show this difference clearly: statically safe variant records typically require the programmer to specify what code will get executed for any possible type. Indeed, this is the best one can hope when the object can indeed vary at run-time. In contrast, cJ statically ensures that the legal operations on an object are fully known.

Configuring generic code is also reminiscent of techniques in C++ template programming [1, 13]. Fundamentally, C++ templates offer a powerful (Turing-complete) but unsafe language for configuring types: there is little static checking capability beyond the checking of templates after instantiation. Furthermore, there is no way to guarantee that a template computation will even terminate. The C++ community has developed ideas on statically validating the input to a template [21, 26] and the general idea of *concepts* has emerged and even developed as a language-independent notion [14]. Nevertheless, concept-based techniques concentrate on validating the type parameters of a generic class, rather than configuring it under static conditions. cJ offers the ability to configure classes based on subtyping conditions, without sacrificing static type safety and with a smooth integration in the base language.

cJ can be viewed as an instance of the aspect-oriented programming paradigm [17], because of its ability to allow a class to be configured based on the structure of a different type hierarchy (representing a cross-cutting concern). As we demonstrated with the cJ implementation of the JCF, the cross-cutting concern “modifiability” is separated in the type system from the intrinsic form of a data structure (e.g., whether it is a list, or a set, or a map). The type system does maintain concepts such as “modifiable list”, “unmodifiable map”, etc., but these are derived from their component types. In fact, the cJ reimplement of the JCF can be compared to previous work that uses AOP to enforce consistency in data structure and behavior [18, 23]—in the JCF, consistency in the usage of data structures along cross-cutting dimensions is enforced by the cJ type system. Nevertheless, cJ differs from common aspect-oriented languages like AspectJ [16] in that separation of concerns in cJ is confined to the type level. cJ does not offer any cross-cutting features at the level of code or method definitions: these are interspersed throughout traditional Java language components (i.e., classes) and not collected in a single entity. Thus, what cJ has to offer is orthogonal to traditional aspect languages and it is interesting to consider integrating their advantages in future work.

8. Conclusions

We presented cJ: an extension of Java that allows declaring class and interface members and supertypes provisionally, under subtyping conditions on parameter types. cJ's power lies in that it allows the composition of orthogonal type hierarchies concisely (avoid-

ing a combinatorial blowup of the number of declared types) yet without sacrificing static type safety. Thus, cJ has a cross-cutting flavor at the level of type hierarchies: the user can define separate aspects of a type hierarchy independently and combine them using cJ type-conditionals to form the complete set of expressible types.

We believe that cJ offers an interesting combination of expressiveness and safety, together with a smooth integration with a representative mainstream OO language. cJ's ability to solve real problems is demonstrated by applying it to the Java Collections Framework. cJ addresses the Collection Framework's well-known shortcomings, eliminating the possibility of run-time errors for unsupported operations without sacrificing conciseness.

Acknowledgments

Phil Wadler and Martin Odersky suggested the need for conditional subtyping and offered many excellent comments that helped strengthen the paper. Multiple anonymous reviewers also made very helpful suggestions. We gratefully acknowledge support by the NSF under Grant CCR-0238289.

References

- [1] A. Alexandrescu. *Modern C++ Design*. Addison-Wesley, 2001.
- [2] G. Bracha and W. Cook. Mixin-based inheritance. In N. Meyrowitz, editor, *Proceedings of the Conference on Object-Oriented Programming: Systems, Languages, and Applications / Proceedings of the European Conference on Object-Oriented Programming*, pages 303–311, Ottawa, Canada, 1990. ACM Press.
- [3] G. Bracha, M. Odersky, D. Stoutamire, and P. Wadler. Making the future safe for the past: Adding genericity to the Java programming language. In C. Chambers, editor, *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, pages 183–200, Vancouver, BC, 1998.
- [4] W.-N. Chin, F. Craciun, S.-C. Khoo, and C. Popeea. A flow-based approach for variant parametric types. In *Proceedings of the 2006 ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages & Applications (OOPSLA '06)*, volume 41, pages 273–290, New York, NY, USA, 2006. ACM Press.
- [5] D. Draheim, C. Luttheroth, and G. Weber. A type system for reflective program generators. In *Generative Programming and Component Engineering (GPCE)*, 2005.
- [6] S. Ducasse, O. Nierstrasz, N. Schärli, R. Wuyts, and A. P. Black. Traits: A mechanism for fine-grained reuse. *ACM Trans. Program. Lang. Syst.*, 28(2):331–388, 2006.
- [7] B. Emir, A. Kennedy, C. Russo, and D. Yu. Variance and generalized constraints for C# generics. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*. Springer, 2006.
- [8] A. Garrido and R. Johnson. Analyzing multiple configurations of a C program. In *21st International Conference of Software Maintenance*. IEEE, 2005.
- [9] S. S. Huang, D. Zook, and Y. Smaragdakis. Statically safe program generation with SafeGen. In *Generative Programming and Component Engineering (GPCE)*, pages 309–326, 2005.
- [10] S. S. Huang, D. Zook, and Y. Smaragdakis. cJ: Enhancing Java with safe type conditions. Technical report, 2006. <http://www.cc.gatech.edu/~ssh/cj/cjfull.pdf>.
- [11] A. Igarashi, B. Pierce, and P. Wadler. Featherweight Java: A minimal core calculus for Java and GJ. In L. Meissner, editor, *Proceedings of the 1999 ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages & Applications (OOPSLA '99)*, volume 34(10), pages 132–146, N. Y., 1999.
- [12] A. Igarashi and M. Viroli. Variant parametric types: A flexible subtyping scheme for generics. *ACM Trans. Program. Lang. Syst.*, 28(5):795–847, 2006.

- [13] ISO Standards Committee. ISO/IEC standard 14882: Programming languages – C++, 1998.
- [14] J. Järvi, J. Willcock, and A. Lumsdaine. Associated types and constraint propagation for mainstream object-oriented generics. In *ACM Symposium on Object Oriented Programming: Systems, Languages, and Applications (OOPSLA)*, pages 1–19, 2005.
- [15] A. Kennedy and D. Syme. Design and implementation of generics for the .NET Common Language Runtime. In *Conference on Programming Language Design and Implementation (PLDI)*. ACM, 2001.
- [16] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. Griswold. Getting started with AspectJ. *Communications of the ACM*, 44(10):59–65, 2001.
- [17] G. Kiczales, J. Lamping, A. Menhdhekar, C. Maeda, C. Lopes, J.-M. Loingtier, and J. Irwin. Aspect-oriented programming. In M. Akşit and S. Matsuoka, editors, *Proceedings European Conference on Object-Oriented Programming*, volume 1241, pages 220–242. Springer-Verlag, Berlin, Heidelberg, and New York, 1997.
- [18] P. Lam, V. Kuncak, and M. Rinard. Crosscutting techniques in program specification and analysis. In *AOSD '05: Proceedings of the 4th international conference on Aspect-oriented software development*, pages 169–180, New York, NY, USA, 2005. ACM Press.
- [19] B. Liskov. *CLU Reference Manual*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 1983.
- [20] V. Litvinov. Constraint-based polymorphism in Cecil: Towards a practical and static type system. In *OOPSLA '98 Conference Proceedings*, volume 33(10), pages 388–411, 1998.
- [21] B. McNamara and Y. Smaragdakis. Static interfaces in C++. In *C++ Template Programming Workshop*, Oct. 2000.
- [22] A. C. Myers, J. A. Bank, and B. Liskov. Parameterized types for Java. In *Conference Record of POPL '97: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 132–145, New York, NY, 1997.
- [23] D. J. Pearce and J. Noble. Relationship aspects. In *AOSD'06: Proceedings of the 5th International Conference on Aspect-oriented Software Development*, pages 75–86, New York, NY, USA, 2006. ACM Press.
- [24] D. Rémy. Type checking records and variants in a natural extension of ML. In *Symposium on Principles of Programming Languages (POPL)*, pages 77–88. ACM Press, 1989.
- [25] N. Scharli, S. Ducasse, O. Nierstrasz, and A. Black. Traits: Composable units of behavior. In *European Conference on Object Oriented Programming (ECOOP)*. Springer LNCS 2743, 2003.
- [26] J. Siek and A. Lumsdaine. Concept checking: Binding parametric polymorphism in C++. In *C++ Template Programming Workshop*, Oct. 2000.
- [27] Y. Smaragdakis and D. Batory. Implementing layered designs with mixin layers. In *Proceedings of the European Conference on Object-Oriented Programming (ECOOP)*, pages 550–570. Springer-Verlag LNCS 1445, 1998.
- [28] M. Torgersen, E. Ernst, and C. P. Hansen. Wild fj. In *FOOL 2005: Foundations of Object-Oriented Languages*, Long Beach, California, 2005. ACM Press.
- [29] M. Torgersen, C. P. Hansen, E. Ernst, P. von der Ahé, G. Bracha, and N. Gafter. Adding wildcards to the Java programming language. In *OOPS track of the Symposium on Applied Computing (SAC)*, 2004.
- [30] A. K. Wright and M. Felleisen. A syntactic approach to type soundness. *Information and Computation*, 115(1):38–94, 1994.
- [31] D. Yu, A. Kennedy, and D. Syme. Formalization of generics for the .NET common language runtime. In *31st Symposium on Principles of Programming Languages (POPL)*. ACM, 2004.