Java Needs Compound Types

Martin Büchi

Turku Centre for Computer Science
Lemminkäisenkatu 14A, FIN-20520 Turku
Martin.Buechi@abo.fi

Wolfgang Weck

Turku Centre for Computer Science
Lemminkäisenkatu 14A, FIN-20520 Turku
Wolfgang.Weck@abo.fi

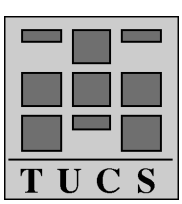

Turku Centre for Computer Science
TUCS Technical Report No 182
May 1998

ISBN 952-12-0224-6
ISSN 1239-1891

Abstract

We propose to extend the Java type system with a new construction: compound types, compositions of a class and any number of interfaces.

Java users will benefit from the type system's increased expressiveness, resulting in more kinds of errors being statically excludable. Compound types are particularly beneficial when several independently predefined standards need to be considered, as typically required by component software. Compound types are a strict extension of Java, i.e., existing programs need not to be modified. Our proposal can be implemented using the existing Java Virtual Machine.

Our scientific contribution is a new combination of name and structural type equivcalence, which our Java extension exemplifies. Type matching is decided by structural comparison of compositions of named types; the latter being only compatible if explicitly declared so. This concept is applicable to several other languages than Java, as well.

We support the feasibility of our proposal by a formal and mechanically verified proof of the extended type system's soundness.

Keywords: component software, combining name and structural equivalence of types, behavioral typing, intersection types, Java, language extension

TUCS Research Group
Programming Methodology Research Group

1 Introduction

One of several reasons to use Java is its support of component-oriented programming, the creation of compiled building blocks to be used in different contexts, and the assembly of systems from such components. JavaBeans [34], Java's component model, competes with other component software standards, such as CORBA [13] and Microsoft's COM [31], but the language itself may also be used to program to these language independent standards.

Type systems, such as Java's, help to document and safeguard component interfaces. By annotating inter-component call parameters with types, one provides some primitive documentation on how to use a service and at the same time expresses a statically checkable precondition: the object passed must implement certain methods, as stated by the type.

Explicitly declared and named types can stand for contracts about services. The behavioral specification is documented separately and linked to the type via the name. A compiler can, of course, not check compliance with such a specification, but it can verify that references to the same types, and intentionally the same contracts, have been made. Explicitly stated contracts are particularly important in the component software realm [35].

Frequently, classes need to conform to more than one contract. For instance, Microsoft's OLE [6] defines ActiveX controls via a bundle of contracts to be implemented. Java supports multiple subtyping to this end. Similarly, one may want to declare variables or method parameters of a type comprising several contracts. This is not supported to the same degree by Java. On a first glance it may seem to be no problem because one only would have to declare the right subtype. We will demonstrate, however, that this may not be possible with independently developed software components.

This problem is explained in Sect. 2. In Sect. 3 we take the problem to its root, the question whether type compatibility is being decided by name equivalence or structural match. We will show that we need a mixture of both, in a way not supported by any language we know. In Sect. 3 we show how to add this to Java by introducing compound types.[1] They are a strict extension, that is, existing Java programs need not be changed. This also holds for the run-time support, the byte code and the virtual machine in particular. We illustrate the latter in Sect. 5. In Sect. 6 we report on a mechanically verified soundness proof for the extended type system. Compound types can be introduced to other typed languages with multiple subtyping as well, as Sect. 7 demonstrates on the example of C++. Section 8 relates to other work and Sect. 9 summarizes our conclusions. In the appendix we collect the proposed changes to Java's language specification.

[1] The compound types defined in this paper are not to be confused with structured types (records, arrays, functions), which are sometimes also called compound types.

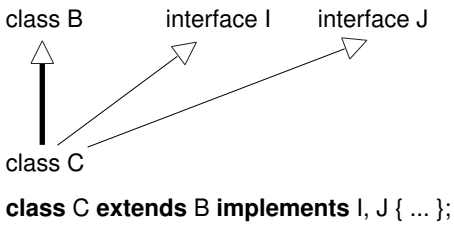

Figure 1: A class extending a base class and implementing two interfaces

2 The Problem

The problem with Java's type system is explained in this section. In 2.1 we briefly review the relevant aspects of Java's type system. In the second subsection, we introduce the essentials of component software, the domain in which the problem mainly surfaces. With these preliminaries we show that it may be impossible to sharply type a parameter to demand a specific combination of interfaces, so that existing or independently developed classes need not be modified to be compatible.

2.1 Java's type system

An essential ingredient of object-oriented programming and component software is polymorphism. Java's type system supports polymorphism through subtyping in two ways: subclassing and implementation of interfaces. Subclassing provides for code inheritance in addition to subtyping, whereas interfaces are pure types, that is, no code is attached to them. Typically, multiple subtyping needs less extra conflict resolution rules than multiple subclassing. Thus, Java's designers decided that a class can have only a single direct superclass but implement several interfaces. Because a class without an explicitly declared superclass implicitly inherits from a predefined class Object, every class —except Object— subtypes exactly one class directly and an arbitrary number, possibly zero, of interfaces, as illustrated in Fig. 1.

If a class C is declared to implement an interface I, all methods being defined by I exist in C, too. Thus, it is type-safe to assign instances of class C to variables of type I.

We have to take a closer look at the situation in which a class C implements several interfaces, say I and J, that both declare a method with the same name, m. If I and J both declare m with exactly the same signature and return type, C defines m only once and binds this to both interfaces. If the parameter lists differ, Java's overloading mechanism takes care of the situation. Both versions exist within C and upon a call the one with the best fitting signature is selected.[2] The case in which the parameter lists are equal but the return types differ is not permitted by

[2]This may not be decidable, in which case an error is flagged, as defined by Java's rules about overloading.

- code inheritance via single subclassing
- multiple subtyping via interfaces (without code inheritance)
- conflict resolution (partially) via overloading
- only types declared to be compatible are compatible (name equivalence)

Figure 2: Relevant aspects of Java's type system

the overloading rules. Consequently, the compiler rejects the declaration of C if I and J conflict in this way.

Subtyping relations can also be established between interfaces. An interfaces I can be declared to extend other interfaces J0, ..., Jn A class implementing I, implicitly implements J0, ..., Jn also, but not vice versa. Interfaces cannot subtype (extend) classes.

Whereas subclassing is basically a mechanism for code reuse, the multiple subtyping offered via interfaces can be used to express different aspects of objects. For instance, a class D of objects being an Applet, runable in a separate thread, and serving as a container in the context of a specific compound document framework, would subclass the class Applet and implement the interfaces Runable and Container.

These combinations are stated without the need to declare a new type first. The class D automatically constitutes such a new type of its own. Another class, E, also extending Applet and implementing the same interfaces as D, establishes a new type too, but a different one. Java uses name equivalence of types, that is two types are compatible only if being declared so. This is not the case for D and E; thus, instances of one cannot be assigned to variables of the other's type. Figure 2 summarizes the discussed aspects of Java's type system.

2.2 Component software

One purpose of using object-oriented technology is to create building blocks to be used in several systems. To support such building blocks, Java features, for instance, separate compilation of classes and bundling of related classes as packages.

Component software tries to move the idea of building blocks to an industrial scale. Like in other engineering disciplines, software systems shall be assembled from premanufactured components rather than crafted individually by hand. This is an old idea, dating back to the NATO conference on software engineering in 1968 [21]. In 1990 Brad Cox even advocated an industrial revolution in the software realm [10], observing that software components are not just a technological issue but a cultural one as well. In particular, as also stressed in a recent book by Clemens Szyperski [35], the true potential of component software comes from establishing component markets. If system assemblers can acquire individual components from several vendors, they can actually combine the many special skills,

ideas, and inventions each of them has to offer. The number of possibly interesting combinations grows rapidly as vendors can join an open market to offer components in the field of their special competence.

An example, frequently used to illustrate this, are spell-checking components that can be composed with a text editor to provide for in-place checking and correction. For instance, a vendor with special competence in linguistics and a specific language, Finnish, offers an add-on spell checker [17] to be used with Microsoft Word. Offering a component rather than a complete word processor allows the company to concentrate on what their staff is good at. Also, the market for Finnish language checkers is probably not extremely large and may not give enough revenue to finance the construction of a competitive editor.

Furthermore, Microsoft's editor serves as an integration platform with other add-ons offered by further vendors. For instance, a user requiring not only well performed checking of Finnish but also needing to include some special type of diagrams into documents, may not find an of-the-shelf program with this particular feature combination. Custom programming of such a system would be too expensive in most cases, but a custom assembly from standard components may be achievable at a reasonable price. Because of this, component software is also described as a path between standard and custom software [28].

Using industrial components as described above has two requirements. Firstly, using a premanufactured component must be easy enough, compared to programming it from scratch, to balance the cost of acquiring it. Otherwise, component reuse is simply not going to happen because system development costs would increase instead of decrease.

The most inexpensive way to compose components is by plug-and-play. This means that neither the components need to be adapted nor any programming is required to glue the components together. The problem may not even be the programming itself, but the required reengineering and detailed understanding of the components or at least their interfaces. Plug-and-play in turn may sometimes be left to the end user. Netscape's Communicator Plug-Ins [7] are an example of this.

As a second requirement, vendors must be able to produce compatible components despite being mutually unaware. In a large component market, it is impossible to demand vendors to know about all other products and to adapt to them.

It seems, as if the latter requirement of independent component development would contradict the former requirement of plug-and-play, but fortunately this is not the case. Component vendors cannot be expected to synchronize their work with each other, but they can build on common standards. If the latter are properly designed, the independently produced components will still interoperate in a plug-and-play manner.

It is not sufficient, however, to use a common wiring (or plumbing) standard, such as COM, CORBA, or a specific programming language, such as Java. These standards only define the calling conventions for procedures or methods respectively. In addition, component plug-and-play requires application domain dependent standard contracts, specified both syntactically and semantically.

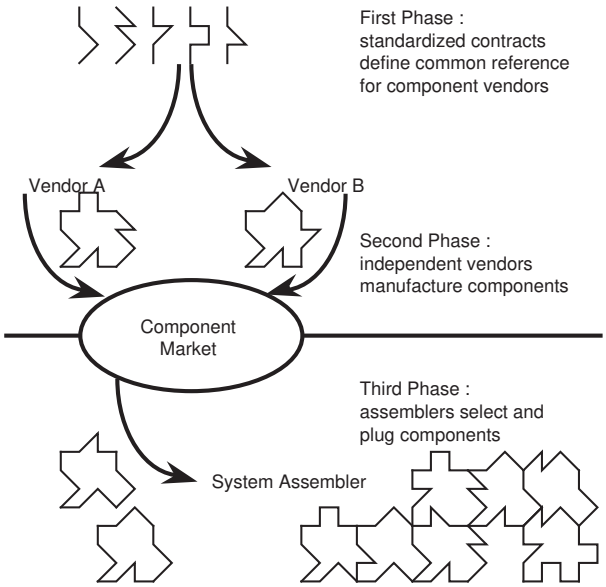

Figure 3: The three phases of component-oriented programming

The three phases of creating component systems are shown by Fig. 3, omitting feedback loops, which drive the evolution of standards and components. During the first phase different standard interfaces are designed and described in public. In the second phase, vendors program towards these standards and place the resulting components on the market. In the third phase, finally, system assemblers select and acquire components from the market and plug them together. Note, that assemblers do not need to analyze *what* a standard interface actually specifies. It suffices to know *that* two components refer to a common standard.

In Java, types are used to support standard contracts. By types we understand both classes and interfaces. Like a plug-and-play system assembler, the loader can check that two components refer to the same type(s) and are thus compatible. A type's name designates the standard. The full behavioral specification associated with it must be stated outside the language in the documentation for component manufacturers.

2.3 A scenario exemplifying a problem with Java

The following example describes a situation that cannot be properly handled by Java's type system. We assume two different and independent standards, which have come into existence entirely unrelated. One of them defines an interface Text, describing operations, such as to insert or delete characters. We also assume a function computing where in a view a specific character would be placed. The second standard defines a compound document framework, like OLE [6], including an interface Container to be implemented by all objects that may act as a compound

document container. The latter must support insertion and removal of document parts. Figure 4 shows portions of these two interfaces in Java.

```
/* as part of a text framework: */
  public interface Text {
    void insert (char ch, int textPos);
    ...
    java.awt.Point displayPoint (int textPos);
    /* returns the display position at which the character at textPos is drawn */
    ...
  };

/* as part of a compound document framework: */
  public interface Container {
    void insertPart (DocPart part, java.awt.Point xyPos);
    ...
  };
```

Figure 4: The standard interfaces Text and Container

Both standards form individually useful frameworks. Vendors can build components for either of them. The problem comes with the wish to create components that build on both standards simultaneously. In our example, this would be components that deal with both texts and containers.

In object-oriented programming, the combination of independently emerged frameworks has been described as an open problem [20]. This is not a problem with component software, where components just implement standard interfaces but do not reuse code from the framework. A component can implement several interfaces belonging to different standard collections. (A common plumbing standard, for instance Java, is still helpful but not strictly necessary.)

```
/* Vendor A's component: */
  public class TextContainerA implements Text, Container {...};

/* Vendor B's component: */
  public class ContainerTextB implements Container, Text {...};
```

Figure 5: Classes offered by vendors A and B

Figure 5 shows portions of two sample classes, TextContainerA of Vendor A and ContainerTextB of vendor B, both implementing the interfaces Text and Container of our sample standards. These classes exhibit a little nuisance. To insert a document part one has to pass the graphical coordinates because the container interface must be used. One may prefer rather to give a text position and have the part inserted after the corresponding character. For this purpose, a generic service can be implemented that maps the text position to a display position and

then inserts the document part there. We assume that a Vendor C wants to offer this service within a class LibraryServices. Figure 6 shows part of this class and Fig. 7 illustrates that whole scenario.

```
public class LibraryServices {
   public static void insertDocPart (DocPart part, ? into, int textPos) { ...};
   /* the question mark stands for a type saying that interfaces
   Text and Container must be implemented */
   ...
};
```

Figure 6: Vendor C's library services

Vendor C's library service works only for instances of classes that implement both interfaces Text and Container. Unfortunately, this cannot be expressed by the type of parameter into, as the question mark in Fig. 6 indicates.

The obvious solution is to create a combined interface TextContainer, which extends both Text and Container and does not add or hide anything, and to declare parameter into of this type. However, neither instances of TextContainerA or ContainerTextB are compatible with the library service, as they are both declared

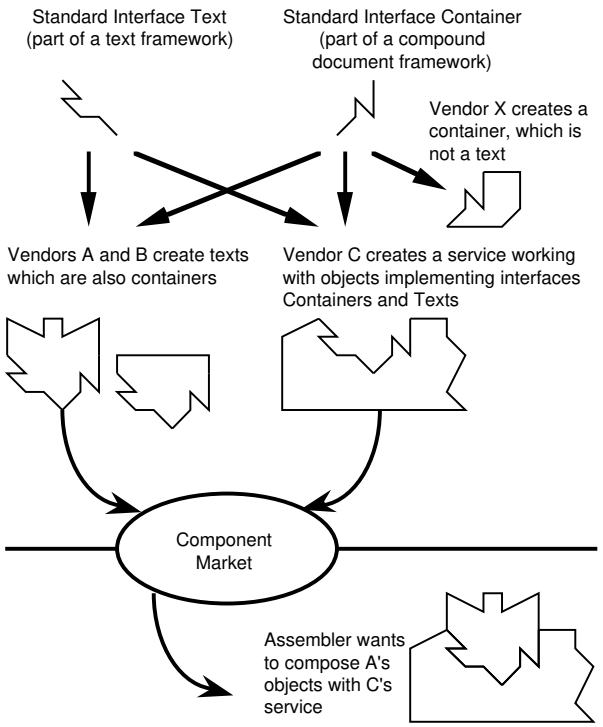

Figure 7: Independent development of classes and insertion service

to implement only the base interfaces but not the combined interface TextContainer. The problem is who is to define the interface TextContainer to be used by all parties? It is not part of either of the two frameworks because they are assumed to be independent. If one of the vendors A, B, or C defines TextContainer, the others would be obliged to use this definition. This contradicts mutual unawareness postulated for component software vendors. On the other hand, if all three vendors declare their own combined interfaces, they are not compatible either.

The problem can be partly tackled by conventions. A class never implements more than one interface directly; an interface never extends any of its superinterfaces by more than one base interface. Instead, combined interfaces named as concatenation of the fully qualified names of the two direct superinterfaces in alphabetical order are introduced into a package CombinedInterfaces. In our example, all three vendors would create and use interface com_X_Text_com_Y_Container, assuming that Text is part of standard X and Container of Y. The system assembler then deletes all but one of the equivalent definitions. Unfortunately, this renders plug-and-play less feasible.

Furthermore, the conventions-based approach suffers from the combinatorial explosion of the number of interfaces. For the combination of three interfaces, the three pair interfaces have to be created and the combined interface has to be defined as the extension of all three pair interfaces as to make its implementations compatible with the pair interfaces as well. The overhead of this solution and, herewith, the pollution of the name space grows exponentially with the number of combined interfaces. Furthermore, legacy classes that do not abide by this convention are left out.

Finally, this approach fails completely, if we try to combine three or more types one of which is a class. Assume that we have a class C and interfaces I and J. According to the above convention, this would give us the abstract classes CI and CJ as well as the combined interface IJ. For classes implementing all three types C, I, and J to be compatible with the three pair types CI, CJ, and IJ, the triple type would have to extend all three pair types. This, however, is impossible because Java does not support multiple class inheritance. We can define a class D which extends CJ and implements IJ, but instances of D —or D's subclasses— cannot be assigned to variables of type CI (Fig. 8). Java does not permit us to declare a class to be compatible with all subsets of its supertypes. This is an additional problem, not bound to component software.

As a different approach, we can resort to run-time tests and textual annotations. We declare parameter into of InsertDocPart (Fig. 6) to be of type Text, add a comment that it must also implement Container, and cast the parameter's value to Container when accessing the latter's members. In this approach, we loose static type checking.

Yet another possibility would be to use two parameters, one of type Text and one of type Container and require them to reference the same object. Again, we need a less desirable run-time test instead of compile-time type checking.

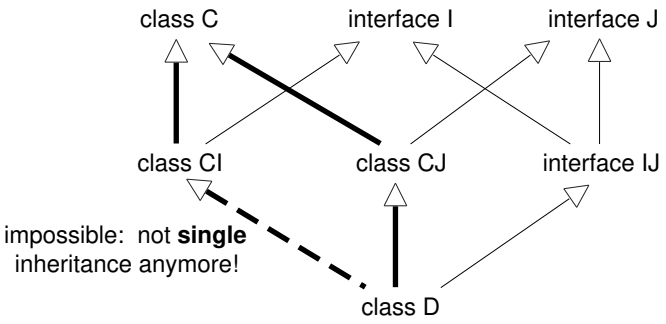

Figure 8: Impossibility of compatibility with all subsets of supertypes

3 Structure vs. Name Equivalence of Types

The problem described above is due to name equivalence of types. The latter can be safely abandoned on the level of type composition, paving the way for compound types.

The fundamental purpose of a type system is to prevent the occurrence of run-time errors [4]. Types prevent, for instance, 'method not understood' and memory access violation errors that could occur if, for example, an integer could be assigned to a pointer. This safety can be achieved already through structural equivalence of types, where any two classes containing methods and fields with the same names and signatures have the same type. Likewise, subtyping is based on structure. The goal is an as big as possible subtyping relation and, therefore, maximal flexibility. Types and the relations between them may even be inferred automatically by the compiler.

If Java would use structural equivalence between types, the problems described in the previous section would not exist. Vendor C could define a combined interface TextContainer, extending Text and Container and use this to type the parameter into. As both implementations TextContainerA and ContainerTextB contain (at least) all methods named in TextContainer, they would be structural subtypes of TextContainer and, thus, compatible with the library service.

The reason against structural type equivalence —and for name equivalence— is that types may be used as an engineering aid for behavioral typing. When doing so accidental matches should be excluded. As an extreme example, borrowed from [19], consider two classes: Rectangle with operations Move and Draw and a class Cowboy with operations Move, Draw, and Shoot. Looking only at the structure, Cowboy is a subtype of Rectangle. To rule this out, Java uses name equivalence between types.

From the modeling perspective, types stand for semantic specifications. Types can be checked by a compiler without the ability to perform semantic proofs. Even if the projections of two unrelated semantic specifications by coincidence result in the same structure, the two should not be considered equal. With name equivalence

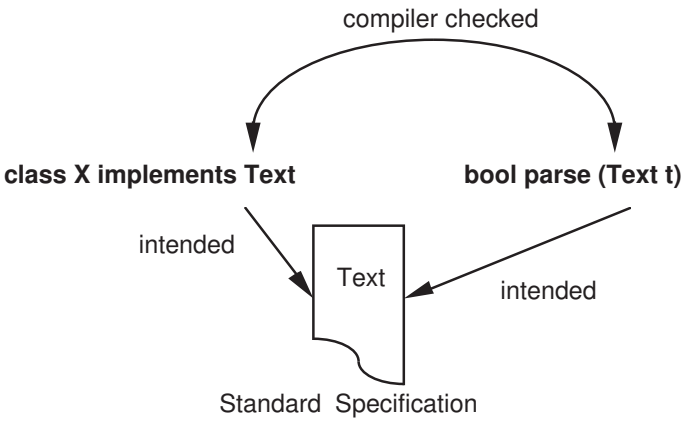

Figure 9: Compiler checked reference to same standard specification

of types, the relations between them must be explicitly declared by the programmer.

Behavioral subtyping [18] and class refinement [24] formalize the idea of behavior associated to types and of subtypes having to refine the behavior of their supertype(s). By simply comparing names, compilers can check that several parties refer to the same standard specification (Fig. 9).

Consider also the analogy to the well-established component market for mass storage with its interfaces standards such as SCSI, IDE, etc. The buyer of a new hard drive simply ensures that she buys a SCSI drive, if that is what it says on her disk controller. For her, the term SCSI represents a common reference made by the manufacturers of the controller and the drive. The buyer is not served well, if she simply shops for a hard disk with matching mechanical connectors, as a drive that adheres to another, incompatible logical signaling standard might also fit this criteria.

Microsoft's COM [31] uses interface identifiers (IIDs) to give each interface its own 'name.' IIDs become, like numbers of ISO standards, abstractions of specifications.

When types are composed from uniquely identified types, structural equivalence suffices also with respect to semantic soundness. If a class like TextContainerA implements both interfaces Text and Container, it must respect both semantic contracts at the same time. One client might, by subsumption, treat an instance of TextContainerA as a Text. By behavioral subtyping, the instance of TextContainerA must behave like a Text, independently of any other interfaces it also implements. Likewise, an instance of TextContainerA can be used as a container. We can also treat an instance of TextContainerA as a Text and as a Container at the same time.

Let us postulate a type CT that requires the classes of its values to implement both interfaces Text and Container. Such a type cannot be expressed in a language

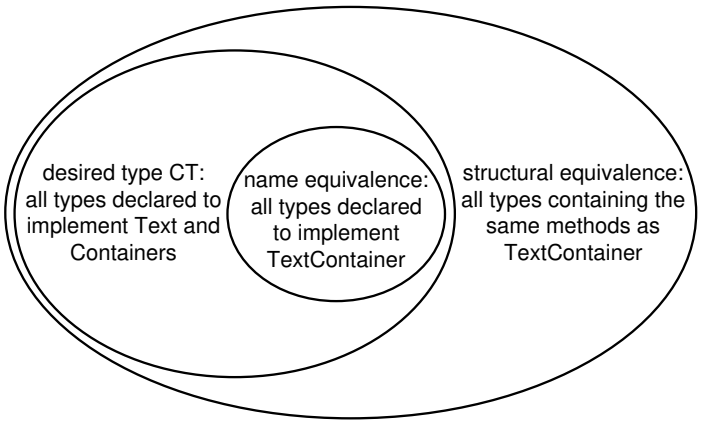

Figure 10: The desired type

based only on structural equivalence. Any class that defines members with the same signature as CT is compatible with the latter, whether it actually implements the two base contracts or not. In a language that supports only name equivalence, the type CT cannot be properly specified either. As we have seen, a combined interface, such as TextContainer, excludes classes that only implement the two base interfaces (Fig. 10).

Using defined types as atoms, the theory of intersection types ([9, 32], see [29] for a recent overview) can be applied. The intersection of two types Σ and Γ is the type of all elements belonging to both Σ and Γ. The new idea is to use the defined types, which represent behavioral contracts, rather than the members of these types as atoms.

To mark this specific choice of atoms and to emphasize the intuition of combining contracts, rather than that of intersecting sets of possible values, we have decided to use the new name *compound type*. The proposed compound types present the required combination of name and structural equivalence to provide maximum flexibility and preserve the semantic notion of typing.

4 Compound Types

In this section we first define compound types and then show how they can be used to solve the typing problems at hand. Finally, we investigate the conditions for well-formedness and a number of interesting properties.

4.1 Definition of compound types

A compound type is an anonymous reference type. It is a direct extension of a set of interfaces and a non-final class, collectively referred to as constituent types. The members (methods, fields) of a compound type are the members of its constituent

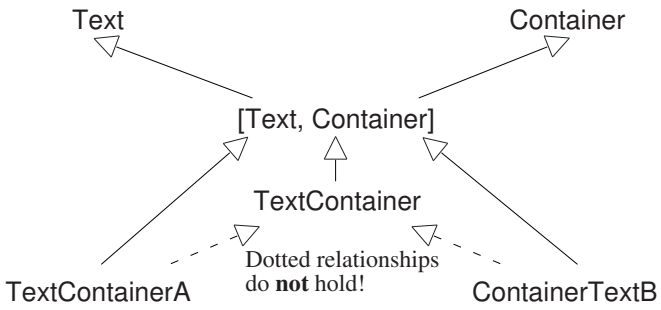

Figure 11: Subtype relationship (transitive arrows omitted)

types with their respective accessibility. The compound type does not add any additional members, redefine any members, or hide any constants. If no constituent class type is explicitly given, Object is implicitly assumed acknowledging that any reference type can be converted to Object by assignment conversion.

Compound types can be used as parameter types, variable types, return types of methods, cast operators, and operands of the instanceof operator. They are not permitted in the **extends** or **implements** clauses of interface and class declarations.

A variable, the declared type of which is a compound type, may have as its value a reference to an instance of a class declared to extend the constituent class and implement the constituent interfaces, or the value of the variable may be null. In other words, the legal values of a variable, the declared type of which is a compound types, are those that could also be assigned to variables of all constituent types of the variables type.

Compound types are written as comma separated lists delimited by square brackets. The order of constituent types is not relevant, e.g., [Text, Container] and [Container, Text] denote the same type.

4.2 Solutions of the typing problems

Using compound types, we can solve our typing problem of the parameter into from Fig. 6. We give it the type [Text, Container]. Since both classes TextContainerA and ContainerTextB implement the two constituent interfaces Text and Container, instances of them can be passed as actual parameters to the library service. Variable into having all members of its constituent interfaces, the required method calls can be made without any additional casts or run-time validity tests. Thanks to the combination of structural and name equivalence, instances of other classes that just happen to declare methods with the same names and signatures as Text and Container, rather than implement the two interfaces, are rightfully rejected at compile time as values for parameter into. Figure 11 illustrates the subtype relationships, omitting transitive arrows.

Compound types also solve the other typing problem pointed out in Sect. 2.3 and illustrated in Fig. 8. Java's type system does not allow a programmer to declare a class that is assignment compatible with all subsets of extended, respectively implemented types, if the class implements more than one interface and does not have Object as its direct superclass. Consider now the case with compound types. Let class G extend class C and implement interfaces I and J. Instances of G can be assigned to variables of types C, I, J, [C, I], [C, J], [I, J], and [C, I, J].

4.3 Well formedness and other properties

As a guiding principle, a compound type [C, I1, I2, ..., In] is well-formed, if the abstract class definition **abstract class** D **extends** C **implements** I1, I2, ..., In {}, where D is a fresh name, would be acceptable in the same package. Thus, no two constituent types may define a method with the same name and signature but different return types.

Including more than one class type is pointless, as no compatible objects could ever be created in Java's single class inheritance system, unless the classes are in a subclass relationship. For simplicity and consistency, we do not allow more than one class to be included.

On the other hand, coherence with Java's design principles dictates that both an interface and one of its superinterfaces may be included. Assume that TrivText is a superinterface of Text. Then, instances of TextContainerA and ContainerTextB can also be assigned to variables of type [Text, Container, TrivText] as their classes indirectly also implement TrivText. However, [Text, Container] and [Text, Container, TrivText] do — due to an in our opinion unfortunate feature of Java — not denote the same type. In Java interfaces may hide constants defined in their superinterfaces. Let TrivText defines a constant **int** k = 21 and Text hide it by defining **boolean** k = **true**. Then a.k denotes the boolean expression **true** for a of type [Text, Container], but is ambiguous for a of type [Text, Container, TrivText]. In the latter case, a qualification such as Text.a.k would be required. Including an interface that is already implemented by the constituent class is analogous. As in a class declaration, the same interface may not be included more than once in a compound type.

Structural equivalence on the level of interface combinations does not introduce any accidental matches. If we declare a variable to be of one of the constituent types, rather than of the compound type, we can always access members of the other types using the corresponding cast. Comparing this to the structural equivalence of classes and interfaces, which cannot be emulated by a cast in Java, reveals why compound types preserve Java's potential of behavioral typing.

Compound types increase type safety compared to today's Java where only one interface can be statically asserted and members of the others are accessed through casts.

We have introduced compound types as anonymous types only. They could, however, also be given names for documentation purposes. Of course, partially

structural equivalence as described above would also apply to the named variant. A named compound type would only be visible where all its constituting interfaces are visible.

5 Emulating Compound Types on the Virtual Machine

Our proposed extension requires modifications of the Java compiler, but programs with compound types can be executed on an unchanged virtual machine [16]. The latter is significant because many of the security and portability properties of Java are tied to the virtual machine, as remarked by Agesen et al. [1].

There are —at least— two ways of emulating compound types. Both of them are sketched in Sect. 2.3. One idea is to use one of the constituent types in place of the compound type and to employ explicit casts to access members of the other constituent types. Most of these casts require run-type validity checks. However, if the byte-code has been generated by a correct compiler all these casts will always succeed at run time, as they have already at compile time been proven correct.

These superfluous run-time tests are also needed when using current Java. If adapting the virtual machine for compound types, an actual performance increase over the state of the art results.

The other way is to use multiple variables, one for each constituent type. These variables all contain the same value but have different types. The compiler takes care that the variables are changed in lockstep; a run-time check that all refer to the same object is not necessary. This solution also comes with some overhead in space and time.

The result of the instanceof expression E **instanceof** [C, L1, L2, …, Ln] is true, at run time, if the value of E is not **null** and the reference could be cast to [C, L1, L2, …, Ln] without raising a ClassCastException. Let o be a fresh variable of type Object. Then

$$\begin{array}{ll} & \mathsf{E}\ \textbf{instanceof}\ [\mathsf{C}, \mathsf{L1}, \mathsf{L2}, \ldots, \mathsf{Ln}] \\ \equiv & (\mathsf{o}{=}\mathsf{E})\ \textbf{instanceof}\ \mathsf{C}\ \&\ \mathsf{o}\ \textbf{instanceof}\ \mathsf{L1}\ \&\ldots\ \&\ \mathsf{o}\ \textbf{instanceof}\ \mathsf{Ln} \end{array}$$

In spite of the emulation option on the existing virtual machine, compound types cannot be mapped to plain Java without abandoning static safety.

6 Type Soundness

In this section, we report on a mechanically verified formal proof of type soundness of Java with compound types. Type soundness intuitively means that all values produced during any program execution respect their static types. An immediate corollary of type soundness is that method calls always execute a suitable method, that is, there are no 'method not understood' errors at run time. Type soundness is not a trivial property, especially for polymorphic languages [3, 4]. It came to prominence with the discovery of its failure in Eiffel [8, 23].

Our proof of type soundness for compound types is based on the work of von Oheimb and Nipkow [36], a much extended version of [25], in which they have formalized and proved the type soundness of a large subset of Java. They verified the proof mechanically with the theorem prover Isabelle/HOL [26].

To this formalization, we added compound types as reference types, appended the widening and casting relations with compound types, and defined the members of the latter. Finally, we adapted the proofs and ran them through Isabelle/HOL.[3] The definition of compound types adds 131 lines to the existing 1371 lines, approximately 10

Here, we present the extensions to the widening and casting relations, which are interesting in their own rights. A full report of all the mechanical details is beyond the scope of this paper.

The Java language specification introduces identity and irreflexive widening conversions separately. Since in all conversion contexts permitting widening identity conversions are possible as well, the two are merged in the formalization. The expression $\Gamma \vdash S \preceq T$ says that in program environment Γ objects of type S can be transformed to type T by identity or widening conversion.[4] In particular, expressions of type S can be assigned to variables of type T and expressions of type S can be passed for formal parameters of type T. Widening can be understood as a syntactic, declared form of subtyping.[5] Unlike subtyping in most type theoretic frameworks, the Java language specification does not say that widening is transitive. Hence, transitivity is a proved property rather than an axiom.

We use the following naming conventions:

C, D	classes	M, L	sets of interfaces
I, J	interfaces	S, T	arbitrary types
R	reference type	Γ	program, environment

Likewise, $\Gamma \vdash C \prec_{\mathsf{C}} D$ expresses that C is a subclass of D, $\Gamma \vdash C \rightsquigarrow I$ that class C implements interface I, and $\Gamma \vdash I \prec_{\mathsf{I}} J$ that I is a subinterface of J. Furthermore, $\mathsf{is_type}\ \Gamma\ T$ expresses that T is a legal type in Γ, $\mathsf{RefT}\ R$ denotes reference type R, and NT stands for the null type. With this, we can express the following two typing judgments, which are also applicable to compound types:

$$\frac{\mathsf{is_type}\ \Gamma\ T}{\Gamma \vdash T \preceq T} \qquad \frac{\mathsf{is_type}\ \Gamma\ (\mathsf{RefT}\ R)}{\Gamma \vdash \mathsf{NT} \preceq \mathsf{RefT}\ R}$$

Further $\mathsf{Class}\ C$ stands for the class type C, $\mathsf{Iface}\ I$ for the interface type I, and $\mathsf{T}[\cdot]$ for an array type with elements of type T. $\mathsf{Compound}\ (C, L)$ denotes the compound type with class C and interfaces $L_i \in L$. The discriminators $\mathsf{is_class}\ \Gamma\ C$,

[3] At http://www.abo.fi/~mbuechi/publications/CompoundTypes.html the Isabelle theories are available.

[4] The turnstile '$\vdash$' stands for the provability of a type relation in a context Γ. As such it is not to be confused with provability of logic formulae, as parts of which type judgments appear.

[5] For simplicity, the term 'subtyping' is used in the other sections of this paper in place of the formally correct notion of 'widening'.

$\Gamma \vdash S \preceq T$	S widens to ('is subtype of') T in Γ
$\Gamma \vdash C \prec_{\mathsf{C}} D$	C is a subclass of D in Γ
$\Gamma \vdash C \rightsquigarrow I$	C implements I in Γ
$\Gamma \vdash I \prec_{\mathsf{I}} J$	I is a subinterface of J in Γ
$\Gamma \vdash S \preceq_{?} T$	cast from S to T permissible at compile time in Γ
'no_conflict Γ (I, D, M)'	in Γ interface I, class D, and constituent interfaces of M do not define a method with the same name and the same signature but different return types

Figure 12: Summary of notation

is_iface Γ I, and is_compound Γ (C, L) are also used. The latter is true, if the compound type is well formed, that is, all constituent types are accessible, C is not final, and there is no method name p such that two constituent types define a method named p with identical signature but different return types. We assume that is_class Γ Object and is_iface Γ Cloneable holds for all Γ. With this we can define the remaining widening rules involving compound types:

$$\frac{\Gamma \vdash \text{Class } C \preceq \text{Class } D;\ \text{is_compound } \Gamma\ (D, M);\ \forall J \in M.\Gamma \vdash C \rightsquigarrow J}{\Gamma \vdash \text{Class } C \preceq \text{Compound } (D, M)}$$

$$\frac{\text{is_iface } \Gamma\ I;\ \text{is_compound } \Gamma\ (\text{Object}, M);\ \forall J \in M.\Gamma \vdash I \prec_{\mathsf{I}} J \vee I = J}{\Gamma \vdash \text{Iface } I \preceq \text{Compound } (\text{Object}, M)}$$

$$\frac{\text{is_type } \Gamma\ T}{\Gamma \vdash \mathsf{T}[\cdot] \preceq \text{Compound } (\text{Object}, \{\})}$$

$$\frac{\text{is_type } \Gamma\ T}{\Gamma \vdash \mathsf{T}[\cdot] \preceq \text{Compound } (\text{Object}, \{\text{Cloneable}\})}$$

$$\frac{\Gamma \vdash \text{Class } C \preceq \text{Class } D;\ \text{is_compound } \Gamma\ (C, L)}{\Gamma \vdash \text{Compound } (C, L) \preceq \text{Class } D}$$

$$\frac{\text{is_compound } \Gamma\ (C, L);\ \Gamma \vdash C \rightsquigarrow J \vee (\exists I \in L.\Gamma \vdash I \prec_{\mathsf{I}} J \vee I = J)}{\Gamma \vdash \text{Compound } (C, L) \preceq \text{Iface } J}$$

$$\frac{\begin{array}{c}\Gamma \vdash \text{Class } C \preceq \text{Class } D;\ \text{is_compound } \Gamma\ (C, L);\\ \text{is_compound } \Gamma\ (D, M);\ \forall J \in M.\Gamma \vdash C \rightsquigarrow J \vee (\exists I \in L.\Gamma \vdash I \prec_{\mathsf{I}} J \vee I = J)\end{array}}{\Gamma \vdash \text{Compound } (C, L) \preceq \text{Compound } (D, M)}$$

The casting relation $\Gamma \vdash S \preceq_{?} T$ states, that a cast from type S to type T is permissible at compile time, that is, the type cast '(T)e', where e is of type S,

might succeed at run-time. If it can be proven to always fail, the compiler can already flag an error.

If $\Gamma \vdash S \preceq T$ holds, the cast can be proven to always succeed. Otherwise, a run-time validity test must be performed to check whether $\Gamma \vdash R \preceq T$ holds for the run-time type R of the cast operand. The following general casting conversions are applicable to compound types as well:

$$\frac{\Gamma \vdash S \preceq T}{\Gamma \vdash S \preceq_? T} \qquad \frac{\Gamma \vdash \mathsf{RefT}\ S \preceq_? \mathsf{RefT}\ T}{\Gamma \vdash (\mathsf{RefT}\ \mathsf{S})[\cdot] \preceq_? (\mathsf{RefT}\ \mathsf{T})[\cdot]}$$

In the rules below, '$\mathsf{no_conflict}\ \Gamma\ (I, D, M)$' means that there is no method name p such that both I and D or one of the interfaces in M declare a method named p with the same signature but different return types.[6] We use this abbreviation freely for different combinations of classes, interfaces, and sets of interfaces to indicate the absence of a method clash in place of the actual predicates, which are lengthy and technical.

$$\frac{\Gamma \vdash D \prec_{\mathsf{C}} C;\ \mathsf{is_compound}\ \Gamma\ (D, M)}{\Gamma \vdash \mathsf{Class}\ C \preceq_? \mathsf{Compound}\ (D, M)}$$

$$\frac{\Gamma \vdash \mathsf{Class}\ C \preceq \mathsf{Class}\ D;\ \mathsf{is_compound}\ \Gamma\ (D, M);\ \neg(\mathsf{is_final}\ \Gamma\ C);\ \text{'}\mathsf{no_conflict}\ \Gamma\ (C, M)\text{'}}{\Gamma \vdash \mathsf{Class}\ C \preceq_? \mathsf{Compound}\ (D, M)}$$

$$\frac{\mathsf{is_iface}\ \Gamma\ I;\ \mathsf{is_compound}\ \Gamma\ (D, M);\ \text{'}\mathsf{no_conflict}\ \Gamma\ (I, D, M)\text{'}}{\Gamma \vdash \mathsf{Iface}\ I \preceq_? \mathsf{Compound}\ (D, M)}$$

$$\frac{\Gamma \vdash D \prec_{\mathsf{C}} C;\ \mathsf{is_compound}\ \Gamma\ (C, L);\ \neg(\mathsf{is_final}\ \Gamma\ D) \vee (\forall I \in L.\Gamma \vdash D \rightsquigarrow I);\ \text{'}\mathsf{no_conflict}\ \Gamma\ (D, L)\text{'}}{\Gamma \vdash \mathsf{Compound}\ (C, L) \preceq_? \mathsf{Class}\ D}$$

$$\frac{\mathsf{is_compound}\ \Gamma\ (C, L);\ \mathsf{is_iface}\ \Gamma\ J;\ \text{'}\mathsf{no_conflict}\ \Gamma\ (C, L, J)\text{'}}{\Gamma \vdash \mathsf{Compound}\ (C, L) \preceq_? \mathsf{Iface}\ J}$$

$$\frac{\mathsf{is_type}\ \Gamma\ T}{\Gamma \vdash \mathsf{Compound}\ (\mathsf{Object}, \{\}) \preceq_? \mathsf{T}[\cdot]}$$

$$\frac{\mathsf{is_type}\ \Gamma\ T}{\Gamma \vdash \mathsf{Compound}\ (\mathsf{Object}, \{\mathsf{Cloneable}\}) \preceq_? \mathsf{T}[\cdot]}$$

$$\frac{\Gamma \vdash D \prec_{\mathsf{C}} C;\ \mathsf{is_compound}\ \Gamma\ (C, L);\ \mathsf{is_compound}\ \Gamma\ (D, M);\ \text{'}\mathsf{no_conflict}\ \Gamma\ (C, L, D, M)\text{'}}{\Gamma \vdash \mathsf{Compound}\ (C, L) \preceq_? \mathsf{Compound}\ (D, M)}$$

[6] In what is believed to be an ommission from the specification [27], Java checks at compile time only for clashes between methods contained in interfaces, but not for clashes between methods contained in classes and interfaces. Opting for maximum static detection of errors, casts involving compound types are defined to check for all kinds of clashes.

$$\frac{\Gamma \vdash \mathsf{Class}\ C \preceq \mathsf{Class}\ D;\ \mathsf{is_compound}\ \Gamma\ (C,L);\ \mathsf{is_compound}\ \Gamma\ (D,M);\ \text{`}\mathsf{no_conflict}\ \Gamma\ (C,L,D,M)\text{'}}{\Gamma \vdash \mathsf{Compound}\ (C,L) \preceq_? \mathsf{Compound}\ (D,M)}$$

Whereas the widening rules can be considered as a particular instantiation of subtyping for intersection types, the rules for casts are believed to be new.

The currently by von Oheimb and Nipkow formalized subset of Java, on which we build, still does not capture all features. Of them final classes, modifiers, class variables, static methods, interface fields, and methods of the class Object would be relevant for compound types.

The main advantages of a mechanized over a paper-and-pencil proof are additional confidence and support for extensions. We would like to stress the second aspect. Not only did the formalization result in a soundness proof, but the proof tool also reminded us of what all needed to be defined about compound types before the desired properties could be established. Most proof scripts worked without modifications. The fact that all theorems were reproved mechanically for the extended language definition conveys more confidence than the typical adaptation of a paper-and-pencil proof with 'this-should-still-hold' handwaving.

7 Compound Types for Other Languages

The idea of compound types is not restricted to Java. It is useful in any other typed language with multiple subtyping and name equivalence of types. The adaption for other languages with single class inheritance and multiple interface subtyping, such as Objective-C [2] where protocols play a similar role as Java's interfaces, is straightforward. In this section, we look at how compound types can be put into languages with multiple class inheritance, such as C++ [33] or Eiffel [22].

C++ does not have interfaces[7] but supports multiple class inheritance. We focus on dynamically allocated objects; the subtle differences to stack objects is beyond the scope of this paper. Hence, compound types are expressed as lists of pointers to class names. Let D and E be two classes. A variable x of declared type [D*, E*] may have as its value a reference to any object that could be cast to both D* and E* in the same context. This definition, which is almost identical to the one for Java, gets a somewhat different meaning in C++'s type system. The main differences stem from non-virtual base classes, protected and private base classes, the different overloading resolution.

Whereas in Java methods with the same name, signature, and return type inherited along different paths are unified, C++ leaves the choice to the programmer whether common base classes should be shared (virtual) or not. Figure 13 shows the sub-object lattices for classes G and I. Objects of type G contain only one single sub object of type E, because the latter is a virtual base class. References to

[7] Microsoft's COM+ SDK C++ [15] compiler introduces the keyword cointerface which on the language level stands for a fully abstract class and which generates a COM interface.

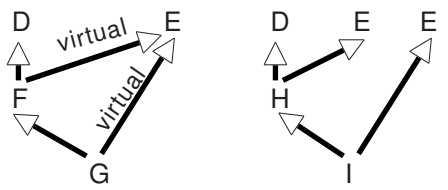

Figure 13: Sub-object lattices of F and I

members of E from objects of type G are unambiguous; hence, references to G can be assigned to x. On the other hand, instances of I contain two distinct sub objects of type E which makes references to them incompatible with [D*, E*].

In C++ base classes can be public, protected, or private. Let K be a subclass of D and E. If, for example, D is a private subclass of K, then objects of type K can only be cast to D within K and its friends. Thus, references to instances of K can be assigned to variables of type [D*, E*] only within K and its friends. Likewise, if D were a protected base class, references to instances thereof can only be assigned to variables of type [D*, E*] within K, subclasses thereof, and friends. This distinction is not necessary in Java because all base classes are public.

If D and E declare methods with the same name, they must be accessed using qualified-ids in classes that inherit from both D and E, even if the methods have different signatures. To be consistent, such methods would also have to be accessed using qualified-ids in variables of type [D*, E*]

Unlike in Java, no compile-time error occurs if two constituent types of a compound type declare a method with the same name and signature but different return types. The methods can be accessed using qualifiers, much like multiple defined constants in Java.

Other features of C++, such as overridding on a per-name base and static methods, also extend to compound types.

8 Related Work

Microsoft's COM: The principle idea of using structural type equivalence with named types as atomic building blocks, each presenting a behavioral contract, is very much inspired by Microsoft's Component Object Model (COM)[31]. In COM, objects cannot be accessed directly but through interfaces only. These interfaces have globally unique identifiers (GUID) as names. It is the intention that with each interface also goes a behavioral specification, to be documented separately.

An object's type is defined as the set of the interfaces implemented by it. This compositional definition of an object's behavior is heavily used, for instance, by the ActiveX Component Framework [6], which defines, for example, an ActiveX control as any object implementing a specific set of interfaces.

Here the parallel ends, however. Clients of a COM object need to use a separate reference variable to each interface through which they want to interact with the object, because each interface may be implemented by a separate node and thus have a different address in memory. This is acceptable because it presents the memory layout under the hood only and can be hidden by a proper programming language. We expect such a language to build heavily on compound typed variables.

To determine whether an object is of a given type, queries must be issued for each interface being part of that type. This may have a serious impact on a system's performance, in particular, if an object is situated remotely. Therefore, categories have been introduced. Membership of classes in a certain category can describe that a specific set of interfaces is supported. Categories can be compared to explicitly declared subtypes in Java. Only if a class makes an explicit reference, that is, registers as a category member, the information can be exploited. No checks are performed to test, whether a class actually implements all the required interfaces of the category it registers in.

Modula-3: Modula-3 [5] is another language which combines name equivalence and structural equivalence of types. This combination, however, is different than what we proposed. In Modula-3 structure equivalence is the default for all types, unless declared as branded, which makes them clearly distinguishable. Also, Modula-3 supports only single subtyping and thus compound types cannot contribute anything.

Intersection types: As pointed out in Sect. 3, intersection types with classes and interfaces as atoms are the theoretical foundation for our approach. Intersection types were introduced into the λ-calculus in the late 70's by Coppo and Dezani-Ciancaglini [9] and independently by Sallé [32]. The original motivation for introducing intersection types was the desire for a type-assignment system in which the typing of terms is invariant under β-expansion and in which every term with a normal form has a meaningfull typing.

In the past twenty years, intersection types, infinite intersections, and the dual notion of union types have been studied extensively in type theory. Pierce and others have also studied the combination of intersection types with bounded polymorphism and other object-oriented concepts (see [29] for a summary of his thesis and an overview of recent work in the field). In contrast to our work, these studies all take the 'type' rather than the 'modeling' view. Thus, they use pure structural equivalence, not taking semantic soundness into account.

Forsythe [30], a descendant of Algol 60, is the only programming language, that uses intersection types and that we are aware of. Forsythe is based on pure structural equivalence, rather than on a combination of name and structural equivalence as our approach. 'Objects' exist in the form of function records only, not allowing for co-variant specialization of the self parameter.

Binary component adaptation (BCA): BCA allows components to be adapted in binary form and during program loading [14]. BCA rewrites class files before or while they are loaded without requiring source code access. Thus, modifications

described by delta files can be applied by third-parties. Adding an interface to the implements clause, one of the supported modifications, could be used to solve the compatibility problem described in Sect. 2.3: Vendor C, the creator of the library service, declares a combined interface, which is used to type the parameter into (Fig. 6). Even if vendors A and B have not declared their components to implement this interface, a component integrator can add it to the lists of implemented interfaces using BCA.

BCA adds further flexibility because it can be used to glue classes that are not based on common standard interfaces. Unfortunately, it also puts the burden of figuring out how to do this correctly on the system assemblers. That is, the system assemblers need to understand the interfaces' semantics and program the adaption. Plug-and-play with made-to-fit components, as enabled by compound types, is the more economical alternative wherever applicable. Furthermore, BCA makes systems harder to understand as delta files must also be taken into account.

BCA does not solve the problem of compatibility with all subsets of supertypes for the case of three or more types including a class other than Object (Fig. 8), because BCA does not add any new kind of types or modify any conversion rules.

Sather: Sather [11], an object-oriented programming language featuring multiple subtyping and subclassing in separate hierarchies, allows the programmer to introduce types as supertypes of already existing classes. This gives in certain cases a partial solution to the compatibility problem described in Sect. 2.3. If vendor C creates the library service after vendor A creates her component and is aware of the latter, then C can declare A's implementation to be a subtype of the combined type introduced by C to type the parameter into (Fig. 6). Unlike BCA, this requires access to C's source code and, therefore, cannot be done a poseriori at system assembly time. Unfortunately, the library service is still not compatible with vendor B's component, as C was not aware of B's implementation and did thus not explicitly declare it to be a subtype. Likewise, any components created after the library service, the manufacturers of which were not aware of the combined type introduced by C, are incompatible with the library service (Fig. 14). With the mutual unawarness postulate for a large component market, Sather's supertyping does, therefore, not solve the problem at hand. Objective-C's [2] mechanism to change the (single) base class of an existing class ('poseAs:') is even less suited for the task at hand as it does not extend to protocols, Objective-C's correspondent to Java's interfaces.

9 Conclusions

We have exhibited a shortcoming of Java's current type system. In a programming language for extensible component software, substitutability of typed objects should neither be decided by the types' name nor just by the structural compatibility of signatures exclusively. Name equivalence, as offered by Java, is too restrictive when composing independently evolved standards or frameworks. Structural

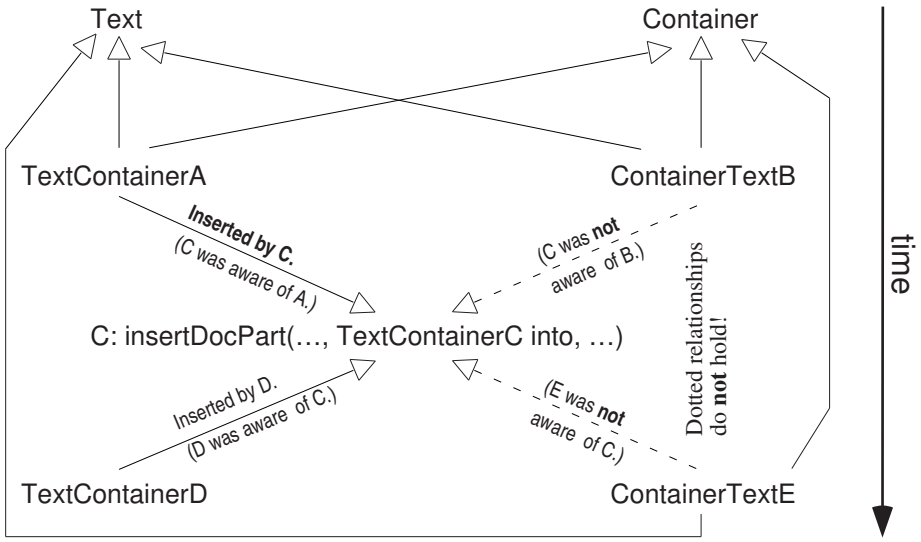

Figure 14: Scenario in Sather where supertyping solves part of problem

equivalence, on the other hand, does not support behavioral typing, that is, to associate semantical specifications with type names.

We concluded that one needs both. On the level of declared types, name equivalence is to be used. A behavioral contract can be associated with each type. When composing these types, however, we want separately declared compositions to be compatible if they have the same structure, that is if they consist of the same types.

This led us to the proposal of compound types to be added to Java. Compound types are anonymous compositions of named types. To a variable of a compound type one can assign any object with the same structure in terms of implemented interfaces and extended classes.

Java is well suited to host compound types. Building on multiple inheritance of interfaces, we integrated our proposal smoothly. The resulting language is a strict extension and thus backward compatible. Java programs with compound types can be executed on an unchanged virtual machine. A mechanical soundness proof gives additional confidence in the well-definedness of the extended type system. The relative ease of adapting a formalization of the existing Java language further illustrates the orthogonality of our proposal.

We believe that compound types can contribute to any typed language with multiple subtyping and name equivalence of types.

Acknowledgments David von Oheimb and Tobias Nipkow provided us with their formalization of Java and helped us with our extensions. We would also like to thank Ralph Back, Dominik Gruntz, Cuno Pfister, and Clemens Szyperski for a number of fruitful discussions.

References

[1] Ole Agesen, Stephen N. Freund, and John C. Mitchell. Adding type parameterization to the Java language. In *Proceedings of OOPSLA '97*, pages 49–65. ACM Press, 1997.

[2] Inc. Apple Computer. Objective-c, 1998. http://developer.apple.com/techpubs/rhapsody/ObjectiveC/index.html.

[3] Kim B. Bruce, Robert van Gent, and Angela Schuett. PolyTOIL: A type-safe polymorphic object-oriented language. In *Proceedings of ECOOP '95*, pages 27–51. LNCS 952, Springer Verlag, 1995.

[4] Luca Cardelli. Type systems. In *Handbook of Computer Science and Engineering*, chapter 103. CRC Press, 1997. http://www.luca.demon.co.uk/Papers.html.

[5] Luca Cardelli, James Donahue, Lucille Glassman, Mick Jordanand Bill Kalsow, and Greg Nelson. Modula-3 report (revised). Research Report 52, Systems Research Center, Digital Equipment Corporation, Palo Alto, November 1989. http://www.research.digital.com/src/m3defn/html/.

[6] David Chappell. *Understanding ActiveX and OLE*. Microsoft Press, 1996.

[7] Netscape Communications. Netscape Plug-Ins (DevEdge Online Documentation), 1998. http://developer.netscape.com/docs/manuals/communicator/ plugin/index.htm.

[8] William Cook. A proposal for making Eiffel type-safe. In *Proceedings of ECOOP '89*, pages 57–70. Cambridge University Press, 1989.

[9] M. Coppo and M. Dezani-Ciancaglini. A new type assignment for λ-terms. *Archiv. Math. Logik*, 19:139–156, 1978.

[10] Brad Cox. Planning the software industrial revolution. *Software Technologies of the 90's special issue of IEEE Software magazine*, November 1990.

[11] B. Gomes, D. Stoutamire, B. Weisssman, and H. Klawitter. Sather 1.1 : Language essentials, 1998. http://www.icsi.berkeley.edu/~sather/Documentation/LanguageDescription/contents.html.

[12] James Gosling, Bill Joy, and Guy Steele. *The Java Language Specification*. Addison Wesley, 1996.

[13] Object Management Group. The common object request broker: Architecture and specification, 1997. Revision 2.0, formal document 97-02-25, http://www.omg.org.

[14] Ralph Keller and Urs Hölzle. Binary component adaptation. In *Proceedings of ECOOP '98*. LNCS, Springer Verlag, 1998. http://www.cs.ucsb.edu/oocsb/papers/ecoop98.html.

[15] Mary Kirtland. The COM+ programming model makes it easy to write components in any language. *Microsoft Systems Journal*, December 1997. http://www.microsoft.com/msj/1297/complus2/complus2.htm.

[16] Tim Lindholm and Frank Yellin. *The Java Virtual Machine Specification*. Addison Wesley, 1996.

[17] Lingsoft. Orthografix: Finnish proofing tools for Microsoft Word, 1998. http://www.lingsoft.fi/.

[18] Barbara H. Liskov and Jeanette M. Wing. A behavioral notion of subtyping. *ACM Transactions on Programming Languages and Systems*, 16(6):1811–1841, November 1994.

[19] Boris Magnusson. Code reuse considered harmful. *Journal of Object-Oriented Programming*, 4(3):8, November 1991.

[20] Michael Mattsson and Jan Bosch. Framework composition: Problems, causes and solutions. In *Proceedings of TOOLS USA 97*, 1997.

[21] McIllroy. Mass-produced software components. In Peter Naur, Brian Randell, and J. N. Buxton, editors, *Software engineering: concepts and techniques: proceedings of the NATO conferences. The Conference on Software Engineering held in Garmisch, Germany, 7th to 11th October 1968*. Petrocelli/Charter, 1976.

[22] Bertrand Meyer. *Eiffel: The Language*. Prentice Hall, second edition, 1992.

[23] Bertrand Meyer. *Object-Oriented Software Construction*. Prentice Hall, second edition, 1997.

[24] Anna Mikhajlova and Emil Sekerinski. Class refinement and interface refinement in object-oriented programs. In *Proceedings of FME'97: Industrial Applications and Strengthened Foundations of Formal Methods*, pages 82–101. LNCS 1313, Springer Verlag, 1997.

[25] Tobias Nipkow and David von Oheimb. Java$_{\ell ight}$ is type-safe — definitely. In *Proc. 25th ACM Symp. Principles of Programming Languages*, pages 161–170. ACM Press, 1998.

[26] Lawrence C. Paulson. *Isabelle: A Generic Theorem Prover*. LNCS 828, Springer Verlag, 1994. see also http://www.cl.cam.ac.uk/Research/HVG/isabelle.html.

[27] Roly Perera and Peter Bertelsen. The unofficial Java spec report, 1997. http://www.nodule.demon.co.uk/java/java-1.0spec-bugs.htm.

[28] Cuno Pfister. Component software: A case study using BlackBox components (online tutorial of the BlackBox Component Builder), 1997. http://www.oberon.ch.

[29] Benjamin C. Pierce. Intersection types and bounded polymorphism. *Mathematical Structures in Computer Science*, 7(2):129–193, April 1997.

[30] John C. Reynolds. Design of the programming language forsythe. In *Algol-like Languages*, volume 1, pages 173–234. Birkhäuser, 1997. Also available as CMU-CS-96-146, ftp://reports.adm.cs.cmu.edu/usr/anon/1996/CMU-CS-96-146.ps.gz.

[31] Dale Rogerson. *Inside COM*. Microsoft Press, 1996.

[32] P. Sallé. Une extension de la théory des types en λ-calcul. In *Proceedings of Automata, Languages and Programming*, pages 398–410. LNCS 61, Springer Verlag, 1978.

[33] Bjarne Stroustrup. *The C++ Programming Language*. Addison Wesley, third edition, 1997.

[34] Sun Microsystems, Inc. Java Beans, 1997. http://splash.javasoft.com/beans/.

[35] Clemens A. Szyperski. *Component Software : Beyond Object-oriented Programming*. Addison-Wesley, 1998.

[36] David von Oheimb and Tobias Nipkow. Machine-checking the Java specification: Proving type-safety. In Jim Alves-Foss, editor, *Formal Syntax and Semantics of Java*. LNCS, Springer Verlag, 1998, to appear.

A Language Manual

This appendix contains the proposed changes to the Java language specification [12]. We describe the extensions of the grammar, the additional conversions, and the members of a compound type with their accessibility. Some other sections of the language specification would also require minor changes, mostly just replacing 'class or interface' by 'class, interface, or compound'. The motivation and definition of compound types are for brevity's sake not repeated. The actual additions are set in italics and surrounded by brief explanations of the context.

A.1 Grammar

The grammar defines the lexical and syntactic structure of a Java program. It is extended as follows, where additions are emphasized:[8]

ReferenceType:
  ClassOrInterfaceType
  CompoundType
  ArrayType
CompoundType
  [CompoundTypeList]
CompoundTypeList
  ClassOrInterfaceType
  CompoundTypeList , ClassOrInterfaceType

A.2 Conversions

Every Java expression has a type that can be deduced from the structure of the expression and the types of the literals, variables, and methods mentioned in the expression. A conversion from type S to type T allows an expression of type S to be treated at compile time as if it had type T instead. In some cases this will require a corresponding action at run time to check the validity of the conversion.

The specific conversions that are possible in Java are grouped into categories. The categories of widening and narrowing conversions are extended for compound types. Additionally, identity and string conversions are also applicable to compound types without changes to the language manual.

There are five conversion contexts in which specific conversions of Java expressions may occur, four of which are relevant for compound types. We refrain from detailing the changes to the conversion contexts, as they are simple instantiations of the specific conversions.

In the rest of this subsection, we describe the specific conversions. We follow the convention that C and D stand for classes and I, J, Lx, and My stand for interfaces.

[8] The extensions apply to the syntactic grammar developed up to Chapter 18 of the language specification. For the LALR(1) grammar in Chapter 19, production CastExpression also requires modification.

Some of the conversion conditions are quite intricate to exclude the identity case from widening and narrowing conversions and also to exclude widening in the narrowing case. The formal widening relations in Sect. 6 of this paper are simpler, because they acknowledge the facts that any context allowing widening conversions also permits identity conversions and likewise casting, the only context that uses narrowing reference conversions also admits identity and widening conversions. For uniformity with the language manual, we have opted for the strict separation; in spite of the fact that the language manual also contains a bug in the case of interface-to-interface narrowing reference conversions.[9]

A.2.1 Widening reference conversions

Widening reference conversions never require a special action at run time and, therefore, never throw an exception. They consist simply in regarding a reference as having some other type in a manner that can be proved correct at compile time. We add the following widening conversions concerning compound types:

1. *From any class type* C *to any compound type* [D, M1, M2, ..., Mm]*, provided that* C *and* D *denote the same class or* C *is a subclass of* D *and provided that* C *implements all constituting interfaces* M1, M2, ..., Mm*.*
2. *From any interface type* I *to any compound type* [Object, M1, M2, ..., Mm]*, provided that for all constituting interfaces* M1, M2, ..., Mm *interface* I *either denotes the same interface or is a subinterface thereof.*
3. *From any array type* A *to compound types* [Object] *and* [Object, Cloneable]*.*
4. *From the null type to any compound type.*
5. *From any compound type* [C, L1, L2, ..., Ln] *to any class type* D*, provided that* C *and* D *denote the same class or* C *is a subclass of* D*.*
6. *From any compound type* [C, L1, L2, ..., Ln] *to any interface type* J*, provided that* C *implements* J*,* J *is one of the constituent interfaces* [L1, L2, ..., Ln]*, or a superinterface of one of them.*
7. *From any compound type* [C, L1, L2, ..., Ln] *to any compound type* [D, M1, M2, ..., Mm]*, provided that* C *and* D *denote the same class or* C *is a subclass of* D *and provided that for all* My *either* C *implements* My *or there exists a* Lx *such that* Lx *is a subinterface of* My *or they denote the same interface.*

A.2.2 Narrowing reference conversions

Such conversions require a test at run time to ensure that the actual reference value is a legitimate value of the new type. If not, then a ClassCastException is thrown.

[9] http://developer.java.sun.com/developer/bugParade/bugs/4133160.html

We add the following narrowing conversions:

1. *From any class type* C *to any compound type ([D, M1, M2, …, Mm]), provided that* C *is a superclass of* D*.*

2. *From any non-final class type* C *to any compound type ([D, M1, M2, …, Mm]), provided that* C *denotes the same class as* D *or* C *is a subclass of* D*, provided that there is no method name* p *such that both* C *and one of the constituting types* D, M1, M2, …, Mm *contain an accessible method named* p *with the same signature but different return types, and provided that* C *does not implement all constituting interfaces* M1, I2, …, Mm*.*

3. *From any interface type* I *to any compound type ([D, M1, M2, …, Mm]), provided that* D *is not* Object *or there exists a* My *such that* I *and* My *do not denote the same interface and* I *is not a subinterface of* My *and provided that there is no method name* p *such that both* I *and one of the constituting types* D, M1, M2, …, Mm *contain an accessible method named* p *with the same signature but different return types.*

4. *From any compound type [C, L1, L2, …, Ln] to any class type* D*, provided that* D *is a subclass of* C*, provided that either* D *is non-final or it implements the constituent interfaces* L1, L2, …, Ln*, and provided that there is no method name* p *such that both* D *and one of the constituting interfaces* L1, L2, …, Ln *contain an accessible method named* p *with the same signature but different return types.*

5. *From any compound type [C, L1, L2, …, Ln] to any interface type* J*, provided that* C *does not implement* J *and none of the constituting interfaces* L1, L2, …, Ln *denotes the same interface as* J *or is a subinterface thereof and provided that there is no method name* p *such that both* J *and one of the constituting types* C, L1, L2, …, Ln *declare an accessible method named* p *with the same signature but different return types.*

6. *From compound types ([Object]) and ([Object, Cloneable]) to any array type.*

7. *From any compound type [C, L1, L2, …, Ln] to any compound type [D, M1, M2, …, Mm], provided that* C *is a superclass of* D *and there is no method name* p *such that constituting types of both compound types declare an accessible method named* p *with the same signature but different return types.*

8. *From any compound type [C, L1, L2, …, Ln] to any compound type [D, M1, M2, …, Mm], provided that* C *denotes the same class as* D *or is a subclass thereof and there is no method name* p *such that constituting types of both compound types declare an accessible method named* p *with the same signature but different return types and provided that there exists a* My *such*

that C *does not implement* My *and there is no* Lx *such that* Lx *denotes the same interface as* My *or is a subinterface thereof.*

A.3 Members of a compound type

The members (fields and methods) of a compound type are all the members of the constituent types. Like classes, compound types may have two or more fields with the same simple name if the fields are declared in different constituting types or supertypes thereof. These fields must be accessed with a qualified name, or a compile-time error occurs. Resolution of overloaded method names works as for classes.

A member of a compound type is accessible only if it is accessible for the constituent type from which it stems.

Turku Centre for Computer Science
Lemminkäisenkatu 14
FIN-20520 Turku
Finland

http://www.tucs.abo.fi

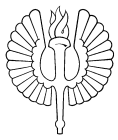

University of Turku
- **Department of Mathematical Sciences**

Åbo Akademi University
- **Department of Computer Science**
- **Institute for Advanced Management Systems Research**

Turku School of Economics and Business Administration
- **Institute of Information Systems Science**