AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Verifying traits: an incremental proof system for fine-grained reuse

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/150050> since 2016-06-29T10:52:54Z

Published version:

DOI:10.1007/s00165-013-0278-3

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

UNIVERSITÀ DEGLI STUDI DI TORINO

This is an author version of the contribution published on:

Ferruccio Damiani, Johan Dovland, Einar Broch Johnsen, Ina Schaefer
Verifying traits: an incremental proof system for fine-grained reuse
FORMAL ASPECTS OF COMPUTING (2014) 26
DOI: 10.1007/s00165-013-0278-3

The definitive version is available at:

<http://link.springer.com/content/pdf/10.1007/s00165-013-0278-3>

Verifying Traits: An Incremental Proof System for Fine-Grained Reuse[†]

Ferruccio Damiani^a and Johan Dovland^b and Einar Broch Johnsen^b and Ina Schaefer^c

^a Università di Torino, Dipartimento di Informatica, Italy

^b Department of Informatics, University of Oslo, Norway

^c Institute for Software Engineering, Technische Universität Braunschweig, Germany

Abstract. Traits have been proposed as a more flexible mechanism than class inheritance for structuring code in object-oriented programming, to achieve fine-grained code reuse. A trait originally developed for one purpose can be adapted and reused in a completely different context. Formalizations of traits have been extensively studied, and implementations of traits have started to appear in programming languages. So far, work on formally establishing properties of trait-based programs has mostly concentrated on type systems. This paper presents the first deductive proof system for a trait-based object-oriented language. If a specification of a trait can be given a priori, covering all actual usage of that trait, our proof system is modular as each trait is analyzed only once. However, imposing such a restriction may in many cases unnecessarily limit traits as a mechanism for flexible code reuse. In order to reflect the flexible reuse potential of traits, our proof system additionally allows new specifications to be added to a trait in an *incremental* way which does not violate established proofs. We formalize and show the soundness of the proof system.

Keywords: Traits, Object Orientation, Program Verification, Proof Systems, Incremental Reasoning

1. Introduction

In object-oriented languages with class inheritance, classes traditionally have three competing roles as generators of objects, as types for objects, and as units of reuse. In contrast, *traits* are pure units of behavior, designed for flexible, fine-grained reuse [SDNB03, DNS⁺06]. A trait contains a set of methods which is completely independent from any class hierarchy. Thus, the common methods of a set of classes can be factored into a trait. Traits can be composed in an arbitrary order. The resulting composite unit (which can be a class or another trait) has complete control over the conflicts which may arise in the composition, and must solve these conflicts explicitly. A trait which was developed for a particular purpose may later be adapted and reused in a completely different context. This way, traits achieve a very attractive level of code reuse, but this flexibility can lead to potentially undesired or conflicting program behavior.

[†] The authors of this paper are listed in alphabetical order. This work has been partially supported by the Deutsche Forschungsgemeinschaft (SCHA 1635/2-1), the Italian MIUR (PRIN 2008 DISCO), the German-Italian University Centre (Vigoni program), and the EU project FP7-231620 HATS: Highly Adaptable and Trustworthy Software using Formal Models (<http://www.hats-project.eu>).

Correspondence and offprint requests to: Einar Broch Johnsen, Department of Informatics, University of Oslo, PO Box 1080 Blindern, N-0316 Oslo, Norway. Email: einarbj@ifi.uio.no.

Since their formulation in SMALLTALK-like languages [SDNB03, DNS⁺06, BDNW08, BDN⁺09], various formulations of traits have been studied for inheritance-based JAVA-like languages (e.g., [SD05, NDS06, RT07, BDG07, BDG08, LS08b, LS08a, BDSS10, BDSS13]). The recent programming language FORTRESS [SAC⁺11] (which has no class-based inheritance) has a trait construct, while the ‘trait’ construct of SCALA [OSV10] is a form of mixin, i.e., a subclass parametrized over its superclass (e.g., [BC90, LM96, FKF98, ALZ03]). Research on ensuring properties of trait-based programs has so far mostly considered type systems (e.g., [RT06, SD05, RT07, BDG08, LS08a, BCD12, BDD⁺10, BDGS13]). These approaches establish that the composed program is type correct; i.e., all required fields and methods are present with the appropriate types.

This paper presents a deductive proof system for verifying behavioral properties of trait-based programs. Flexible, fine-grained code reuse motivates the development of traits as a mechanism for code structuring, but this flexibility poses challenges when developing a corresponding proof system: a high degree of code reuse in a code structuring mechanism typically limits the degree of reasoning reuse which can be supported by the proof system. Ideally, when traits are composed into a class, the trait specifications should already provide enough information to ensure that the interface contracts of that class hold. This scenario would result in a purely compositional proof system in which the actual usage of a trait always corresponds to its originally intended usage, as captured by its original specification. These specifications can be established by modular reasoning; i.e., it is sufficient to analyze each trait once during the verification process. In practice, a priori imposing a fixed specification on a trait may overly restrict the flexibility of code reuse in that trait, which goes against the original motivation for traits.

The challenge in developing a deductive proof system for trait-based programs is to support the *flexibility* offered by traits as a code reuse mechanism, while providing an *incremental* and *compositional* reasoning system. Without restricting the flexibility of traits as a code reuse mechanism, a trait cannot be fully verified independently of its context of composition. In order to align our proof system with this flexibility, traits will be associated with *sets of possible specifications*, and the applicable specifications of a trait depend on its context of composition. By incremental reasoning, new specifications may be added *incrementally* to a trait without violating previous specifications and proofs. When traits are composed, the specifications of the composed traits are selected from the compatible specifications of the constituent traits. Hence, our proof system subsumes modularity and supports modular reasoning for traits when applicable, but extends this modularity to incremental reasoning when required for flexible trait reuse. In particular, the incremental approach supports a gradual extension of existing trait libraries with specifications, driven by the verification of programs using the library.

In order to focus the paper on the particular challenges proposed by trait-based programs, the proof system is presented for a kernel of TRAITRECORDJ [BDSS13, Tra11], a trait-based JAVA-like language with a prototype implementation. The proof system presented in this paper can be used to guarantee that programs obtained through the flexible adaptation and composition of traits satisfy critical requirements, by reasoning modularly and incrementally about traits, trait adaptation, and composition. The analysis of trait-based programs is formalized as an inference system which tracks specification sets for traits, when traits are modified and composed. This inference system adapts previous work on lazy behavioral subtyping [DJOS10], which developed an incremental inference system for late bound method calls by separating the required and provided behavior of methods, to trait modification and composition. The inference system presented in this paper does not depend on a particular program logic. For simplicity, we use a Hoare-style notation to specify the pre- and postconditions of method definitions in terms of proof outlines and do not consider, e.g., class or trait invariants. The main approach of our proof system has previously been presented in a short paper at FTfJP 2011 [DDJS11]. This paper extends [DDJS11] with more examples and explanations, details of the formalization, and the proofs of the main results. As far as we know, no other deductive proof system for trait-based languages has been proposed so far.

The paper is structured as follows. Section 2 gives a high-level introduction to our approach to verifying trait-based programs as an incremental reasoning process. Section 3 introduces the kernel trait language used in this paper and Section 4 presents a specification notation for traits. Sections 5 and 6 introduce the two aspects of the proposed proof system for the considered kernel trait language. Section 5 discusses how the proof system is used to verify basic and composed trait expressions and Section 6 explains how to verify interface contracts for classes. The formal inference system for program analysis is defined in Section 7, which also shows soundness for the proof system. Section 8 illustrates by example how the proof system may be used. Related work is discussed in Section 9. Section 10 concludes the paper and discusses future work. The paper includes three appendixes. Appendix A contains technical definitions omitted in the main paper, Appendix B contains the proofs of the soundness theorem and Appendix C details the verification of the example from Section 8.

Fig. 1. The workflow for verifying trait-based programs, starting by the external requirements imposed on classes by their behavioral interfaces.

2. Verifying Trait-based Programs as an Incremental Reasoning Process

The main goal of the verification process for trait-based programs is to show that the classes that are built from the traits implement the contracts of their declared interfaces. This section gives an informal introduction to this verification process, which consists of two parts: trait analysis and class analysis. Section 7 shows that the overall verification process meets its goal through the soundness of the formalized proof system (Theorem 7.4).

Figure 1 illustrates the envisaged workflow for verifying trait-based programs using our verification process. We assume given a library of traits, annotated with specifications (possibly without any non-trivial specifications). These trait specifications are stored in the *trait environment*. The specifications of a trait contain a pre-/post-condition pair for each method, denoting the method’s guarantee. This guarantee holds if the other methods called by this method satisfy certain properties which are contained as requirements in the method specification. While requirements to *external calls* can be checked via the behavioral interfaces, the guarantees of internal calls depend on how the traits are assembled into classes; *internal calls* are checked with respect to the class environment. During the trait analysis, the specifications of methods in composed trait expressions are derived from the specifications of methods in the basic traits, depending on trait modifiers and composition. The specifications of the methods in the basic traits have to be verified by a suitable proof system before they are stored in the trait environment during the trait analysis. In this setting, the trait environment plays the role of a cache of already established specifications that can be reused if applicable during the analysis of classes. For a trait library, the trait environment can be modularly reused and incrementally extended during the analysis of different programs using the library.

During the class analysis, the verification of a program comprising a set of classes built from traits is driven by the *external requirements* to the classes of the program. These external requirements are specified in the behavioral interfaces which are implemented by the classes of the program. The external requirements to the methods of a class are decomposed into requirements on methods in trait expressions used to build the class. The *class environment* stores the information about the method specifications used to verify the trait expressions in order to retrieve those when verifying the requirements of other methods. When an external requirement is imposed on a method in a trait expression during the class analysis, there are three cases:

1. *Modular verification*: The specifications already contained in the trait environment suffice to prove the external requirements imposed on the method. We obtain *specification reuse* by simply showing that the specifications stored in the trait environment imply the external requirement.
2. *Incremental verification*: The specifications already contained in the trait environment do not suffice to prove the external requirements. The user can add new specifications to the traits, which need to be verified. This verification triggers new trait analysis operations in which external calls are verified based on behavioral interfaces and internal calls based on methods on other traits. During this verification process, we can return to the modular case in which

specifications follow from previous specifications or remain in the incremental case which can repeatedly trigger new trait analysis operations until the analysis is complete.

3. *Non-verifiable case*: The verification process for the program fails if we cannot find provable specifications to be added to the trait environment which resolve the external requirements on the classes.

The approach allows program verification to interact with a trait library without assuming that the library already contains specifications: When verifying the first program built from a set of traits, we can start with an empty trait environment. The verification process will then enter the incremental case and derive appropriate annotations for the methods as required by the program and store them in the trait environment. When another program is built by reusing already verified traits, we can have the modular case such that we are able to reuse the previously established specifications or again enter the incremental case and derive new specifications. In this way, while verifying a set of programs, we stepwise obtain all necessary specifications for the traits that are contained. In practice, we expect that programs will often use a trait in the same way so the potential for specification reuse is good. However, the proposed approach also caters for programs that use some traits in innovative ways, supporting fine-grained specification reuse.

3. A Kernel Language for Trait-Based Programs

For the purposes of this paper we consider a kernel of the trait-based programming language TRAITRECORDJ [BDSS13, Tra11] which highlights the specific features of trait-based programming from a reasoning perspective. In particular, TRAITRECORDJ completely separates the three traditional roles of classes as types, units of code reuse, and object generators. Interfaces and primitive types are the only source language types; fields and methods are typed by source language types, and subtyping reflects the extends-relation in interface declarations. Traits only play the role of units of code reuse and are not types. Class-based inheritance is not present, so classes only play the role of object generators.

A trait in TRAITRECORDJ consists of *provided methods*, which are the methods defined in the trait; *required methods*, which parametrize the behavior of auxiliary methods assumed to be available in a class using the trait; and *required fields*, which are similarly assumed to be available in a class using the trait. The required fields of a trait can be directly accessed in the body of the provided methods of the trait, and the required methods can be called internally in the trait's provided methods. Traits are building blocks which can be used to compose classes and other traits by means of a trait summation operation and a suite of trait alteration operations. Since traits do not introduce any state, a class assembled from traits has to provide the required fields of its constituent traits. The kernel language represents a “normal form” syntax which simplifies the analysis proposed in this paper, since it ensures that trait summation happens as late as possible in a trait composition.

3.1. Kernel Language Syntax

The syntax of the kernel language is given in Figure 2. A program P consists of (lists of) interface declarations ID , trait declarations TD , and class declarations CD . Following [IPW01] we use the overline notation for (possibly empty) sequences. In TRAITRECORDJ, interface declarations ID , method headers H , field declarations F , and method declarations M have a similar syntax as in JAVA (but ignoring, e.g., visibility modifiers and exceptions). For simplicity, in the kernel language the syntax of methods is restricted. For instance, a method may only contain a single **return** statement which must be the last statement in the body of the method, and method parameters are not assignable. These language features are well-understood from the literature (e.g., [Hoa69, Apt81, AdBO09]) and do not impose any new difficulties for our proof system. In the sequel, we will use the notation $body(M)$ to denote the body of a method declaration M . We syntactically distinguish *internal* method calls $m(\bar{e})$ from *external* method calls $v.m(\bar{e})$ (where the variable v is typed by an interface).

A trait declaration TD binds a trait name T to a trait expression. Traits are either basic or composed. A *basic trait* binds a basic trait name T_b to a basic trait expression BTE , whereas a *composed trait* binds a composed trait name T_c to a composed trait expression CTE . Only basic traits can be adapted by means of alteration operations, which explains the normal form restriction of the kernel language, and assembled into composed traits.

A *basic trait expression* BTE , written $\{\bar{F}; \bar{H}; \bar{M}\}$, provides the methods $\bar{M}$ and declares the types of the *required* fields $\bar{F}$ and methods $\bar{H}$ (which can be directly accessed in the bodies of the methods $\bar{M}$). A *trait alteration expression* TAE applies trait alteration operations ao to a basic trait T_b . The operation **exclude** forms a new trait by removing a method from an existing trait. The operation **aliasAs** forms a new trait by giving a new name to an existing method; in particular, when a recursive method is aliased, its recursive invocation refers to the original method. The operation

P	$::= \overline{ID} \overline{TD} \overline{CD}$	<i>program</i>
ID	$::= \text{interface } I \text{ extends } \overline{I} \{ \overline{H}; \}$	<i>interface declaration</i>
H	$::= [S \text{void}]_m (\overline{S} \overline{x})$	<i>method header</i>
S	$::= I \text{boolean} \text{int}$	<i>source language type</i>
F	$::= S f$	<i>field</i>
M	$::= H \{ \overline{S} \overline{x}; \overline{s}; [\text{return } e;] \}$	<i>method</i>
s	$::= v = rhs m(\overline{e}) v.m(\overline{e}) \text{if}(e) \{ \overline{s}; \}$	<i>statement</i>
v	$::= f x$	<i>variable</i>
rhs	$::= m(\overline{e}) e.m(\overline{e}) \text{new } C(\overline{e}) e$	<i>expressions with side effects</i>
e	$::= v \text{this} \text{null} \text{true} \text{false} \dots$	<i>expressions without side effects</i>
TD	$::= \text{trait } T_b \text{ is } BTE$	<i>basic trait declaration</i>
	$ \text{trait } T_c \text{ is } CTE$	<i>composed trait declaration</i>
T	$::= T_b T_c$	<i>trait names</i>
BTE	$::= \{ \overline{F}; \overline{H}; \overline{M} \}$	<i>basic trait expression</i>
CTE	$::= TAE CTE + CTE$	<i>composed trait expression</i>
TAE	$::= T_b \overline{a\overline{o}}$	<i>trait alteration expression</i>
ao	$::= [\text{exclude } m] [m \text{ aliasAs } m]$ $ [f \text{ renameTo } f] [m \text{ renameTo } m]$	<i>trait alteration operation</i>
CD	$::= \text{class } C \text{ implements } \overline{I} \text{ by } \{ \overline{F}; \} \text{ and } CTE$	<i>class declaration</i>

Fig. 2. Kernel language syntax. The following naming convention is used in the syntax: $I \in \text{interface names}$; $T_b \in \text{basic trait names}$; $T_c \in \text{composed trait names}$; $C \in Cid$ (the class names); $f \in \text{field names}$; $m \in Mid$ (the method names); $x \in \text{method parameter or local variable names}$. We further denote by Mid the set of method declarations and by $Label$ the set of field and method names.

renameTo creates a new trait by renaming all occurrences of the name of a required field name or of a required or provided method name from an existing trait. A *composed trait expression* CTE is either a trait alteration expression TAE or the sum of two composed trait expressions. The *symmetric sum* $+$ merges two traits to form a new trait and requires that the summed traits are disjoint; i.e., the summed traits cannot provide identically named methods but they may require identically named (and typed) fields or methods.

Classes in TRAITRECORDJ are assembled from a trait expression by providing the fields required by that trait expression. Thus, a class declaration CD binds a class name C to a set of interfaces $\overline{I}$, which specify the possible types for an instance of the class, fields $\overline{F}$, and a trait expression CTE . All fields are private. For simplicity, class constructors are omitted in the paper; each class C is assumed to have an implicit constructor that behaves like the JAVA constructor $C(\overline{S} \overline{x}) \{ \overline{f} = \overline{x} \}$, where $\overline{S} \overline{f}$ are all the fields of the class.

Example 3.1. As a running example in this paper, we consider the implementation of a bank account. The following trait T_{Acc} provides the basic operations for inserting and withdrawing money:

```

trait TAcc is {                               // no required fields
  boolean validate(int a);                       // required method
  void update(int y);                             // required method
  void deposit(int x) {update(x);}
  void withdraw(int id, int x){
    boolean v = validate(id);
    if (v) {update(-x);}
  }
}

```

A basic account CA_{Acc} may then be defined as follows, where the additional trait TA_{Aux} defines the auxiliary methods required by T_{Acc} :

```

interface IAcc {
  void deposit(int x);
}

```

```

    void withdraw(int id, int x);
}

trait TAux is {
    int bal; int owner;          // required fields
    void update(int y) {bal = bal + y;}
    boolean validate(int id) {return (id == owner);}
}

class CAcc implements IAcc by {int bal; int owner;} and TAcc + TAux

```

Since traits define flexible units for code reuse, TAcc may be combined with other traits to define different account behavior. For example, the class CFeeAcc charges an additional fee whenever the balance is reduced. We define another interface IFeeAccount in order to be able to associate different behavioral specifications later. The class CFeeAcc is composed from the traits TAcc, TFee, and TAux where the method update is renamed to basicUpd.

```

interface IFeeAcc {
    void deposit(int x);
    void withdraw(int id, int x);
}

trait TFee is {
    int fee; int bal;           // required fields
    void basicUpd(int y);       // required method
    void update(int y) {
        basicUpd(y); if (y<0) {bal = bal - fee;}
    }
}

class CFeeAcc implements IFeeAcc
    by {int bal; int owner; int fee;}
    and TAcc + TFee + (TAux[update renameTo basicUpd])

```

The public methods of an object are those listed in the interfaces implemented by its class; the other methods and fields are private to the object. For instance, the only public members of classes CAcc and CFeeAcc are the methods deposit and withdraw.

3.2. Kernel Language Semantics

The semantics of a class composed from traits is defined through the *flattening principle* [DNS⁺06] (see also [NDS06, LSZ09]), which states that the semantics of a method introduced in a class through a trait should be identical to the semantics of the same method defined directly within a class. A flattening function defines the semantics of TRAITRECORDJ by translating TRAITRECORDJ class declarations to JAVA class declarations, and a trait expression to a sequence of method declarations. TRAITRECORDJ interfaces are literally JAVA interfaces and need no translation.

Let a TRAITRECORDJ program be represented by an interface table IT, a trait table TT, and a class table CT, which map interface, trait, and class names to interface, trait, and class declarations, respectively. For simplicity, we assume fixed, global tables IT, TT, and CT. The flattening function $\llbracket \cdot \rrbracket$, defined in Figure 3, maps a TRAITRECORDJ class declaration to a JAVA class declaration, and a trait expression to a sequence of method declarations. (The implicit class constructor of TRAITRECORDJ is unaltered by the flattening function, and omitted in Figure 3.) We denote by $\llbracket CD \rrbracket$ the flattened version of a class declaration CD, and by $\llbracket TE \rrbracket$ the method declarations which result from flattening the (basic or composed) trait expression TE.

The flattening function uses some auxiliary functions, which are defined in Figure 4. Given a set $\bar{M}$ of method definitions, let $m \in \bar{M}$ denote that m is defined in $\bar{M}$, and let $\bar{M}(m)$ return the definition of m in $\bar{M}$. Function $rem : Set[Mtd] \times Mid \rightarrow Set[Mtd]$ takes a set of method definitions and a method name m , and returns the set without the definition of m . Function $ren : Mtd \times Mid \times Mid \rightarrow Set[Mtd]$ performs method renaming. The function $rep : Set[Mtd] \times Label \times Label \rightarrow Set[Mtd]$ is defined such that:

$$\begin{aligned}
\llbracket \text{class } C \text{ implements } \bar{T} \text{ by } \{ \bar{F}; \} \text{ and } CTE \rrbracket &\triangleq \text{class } C \text{ implements } \bar{T} \{ \bar{F}; \llbracket CTE \rrbracket \} \\
\llbracket \{ \bar{F}; \bar{H}; \bar{M} \} \rrbracket &\triangleq \bar{M} \\
\llbracket \bar{T} \rrbracket &\triangleq \llbracket TE \rrbracket \quad \text{if } TT(T) = \text{trait } T \text{ is } TE \\
\llbracket CTE_1 + CTE_2 \rrbracket &\triangleq \llbracket CTE_1 \rrbracket \cdot \llbracket CTE_2 \rrbracket \\
\llbracket TAE[\text{exclude } m] \rrbracket &\triangleq \text{rem}(\llbracket TAE \rrbracket, m) \\
\llbracket TAE[m \text{ aliasAs } m'] \rrbracket &\triangleq \begin{cases} \llbracket TAE \rrbracket \cdot \text{ren}(\llbracket TAE \rrbracket(m), m, m') & \text{if } m \in \llbracket TAE \rrbracket \\ \llbracket TAE \rrbracket & \text{otherwise} \end{cases} \\
\llbracket TAE[f \text{ renameTo } f'] \rrbracket &\triangleq \text{rep}(\llbracket TAE \rrbracket, f, f') \\
\llbracket TAE[m \text{ renameTo } m'] \rrbracket &\triangleq \text{rep}(\llbracket TAE \rrbracket, m, m')
\end{aligned}$$

Fig. 3. Flattening TRAITRECORDJ to JAVA (where TE denotes either a basic trait expression BTE or a composed trait expression CTE).

$$\begin{aligned}
\text{rem}(\emptyset, m) &\triangleq \emptyset \\
\text{rem}(\text{I n}(\bar{T} \bar{x})\{t\} \bar{M}, m) &\triangleq \begin{cases} \bar{M} & \text{if } m = n \\ \text{I n}(\bar{T} \bar{x})\{t\} \text{rem}(\bar{M}, m) & \text{otherwise} \end{cases} \\
\text{ren}(\text{I n}(\bar{T} \bar{x})\{t\}, m, m') &\triangleq \text{I n}[m'/m](\bar{T} \bar{x})\{t\} \\
\text{rep}(\emptyset, l, l') &\triangleq \emptyset \\
\text{rep}(\text{I n}(\bar{T} \bar{x})\{t\} \bar{M}, l, l') &\triangleq \text{I n}(\bar{T} \bar{x})\{t[l'/l](\bar{T} \bar{x})\} \text{rep}(\bar{M}, l, l')
\end{aligned}$$

Fig. 4. Definitions of auxiliary flattening functions, where: l and l' range over field and method names; $n[l'/l]$ denotes l' if $n = l$ and l otherwise; and $t[l'/l]$ denotes the substitution of all occurrences of field l in t by l' and the substitution of all internal calls of method l in t by internal calls of method l' .

- if l is a field name, $\text{rep}(\bar{M}, l, l')$ replaces all occurrences of l in the bodies of the methods $\bar{M}$ by l' ; and
- if l is a method name, $\text{rep}(\bar{M}, l, l')$ replaces all occurrences of l in the headers of the methods $\bar{M}$ by l' and replaces all the internal calls of l in the bodies of the methods $\bar{M}$ by internal calls of l' .

Example 3.2. The flattening of the class CFeeAcc, which is introduced in Example 3.1 above, is as follows (the implicit TRAITRECORDJ constructor is omitted).

```

class CFeeAcc implements IFeeAcc {
  int bal; int owner; int fee;
  void deposit(int x) {update(x);}
  void withdraw(int id, int x){
    boolean v = validate(id); if (v) {update(-x);}
  }
  void basicUpd(int y) {bal = bal + y;}
  boolean validate(int id) {return (id == owner);}
  void update(int y) {
    basicUpd(y); if (y<0) {bal = bal - fee;}
  }
}

```

Example 3.3. This example illustrates the difference between the **aliasAs** and **rename** operations on traits. Consider the following trait TFactorial defining a method factorial that computes the factorial of a natural number by means of an auxiliary method fAux:

```

trait TFactorial =
{ int factorial(int n) { int r = -1; if (n >= 0) r = fAux(n); return r; }
  int fAux(int n) { int r = 1; if (n >= 2) r = n * fAux(n-1); return r; } }

```

The flattening of $\llbracket TFactorial[fAux \text{ aliasAs } fact] \rrbracket$ is the following sequence of three methods (the first two methods are exactly those provided by the trait TFactorial):

```

int factorial(int n) { int r = -1; if (n >= 0) r = fAux(n); return r; }
int fAux(int n) { int r = 1; if (n >= 2) r = n * fAux(n-1); return r; }
int fact(int n) { int r = 1; if (n >= 2) r = n * fAux(n-1); return r; }

```

while the flattening of $\llbracket \text{TFactorial}[\text{fAux renameTo fact}] \rrbracket$ is the following sequence of two methods:

```

int factorial(int n) { int r = -1; if (n >= 0) r = fact(n); return r; }
int fact(int n) { int r = 1; if (n >= 2) r = n * fact(n-1); return r; }

```

The TRAITRECORDJ type system ensures that the flattening of a well-typed TRAITRECORDJ program is a well-typed JAVA program. It has been formalized for FEATHERWEIGHT RECORD-TRAIT JAVA (FRTJ) [BDS10, BDS09], a minimal core calculus for TRAITRECORDJ in the spirit of FEATHERWEIGHT JAVA [IPW01]. A type system for the kernel language considered in this paper would be a straightforward adaptation of the type system for FRTJ; in the sequel, it is assumed that programs are well-typed according to such a type system.

4. Specifying Trait-Based Programs

This section considers the specification of trait-based programs. The proof system developed in this paper does not depend on a particular program logic. Let PL be a (sound) program logic and let a, p, q range over assertions in the assertion language of PL . For simplicity in the presentation we use Hoare triples $\{p\}s\{q\}$ [Hoa69], with a standard partial correctness semantics [Apt81, AdBO09], adapted to the object-oriented setting; in particular, de Boer's technique using sequences in assertions addresses the issue of object creation [dB99]. Thus, the triple $\{p\}s\{q\}$ expresses that if s is executed in a state where the precondition p holds and the execution terminates, then the postcondition q holds after s has terminated. Let $A \vdash_{PL} \{p\}s\{q\}$ denote that $\{p\}s\{q\}$ is derivable from a (possibly empty) set of axioms A using the inference rules of PL . An *assertion pair* (p, q) is a pair of assertions such that p is a precondition and q a postcondition (for some sequence of program statements).

Entailment. The standard rule of consequence in Hoare Logic (e.g., [AdBO09]) is insufficient for dealing with sets of assertion pairs, which we will need to flexibly combine information from assertion pairs. We take a relational approach to entailment, and let q^o denote the assertion q in which all occurrences of fields f and method parameters x have been substituted by the corresponding fields f^o and local variables x^o , respectively, avoiding name capture. The assertion pair (p, q) is understood as an input/output relation $\forall \bar{z} . p \Rightarrow q^o$, where the formula q^o ensures that the postcondition applies to the values of fields and local variables in the post state, $\bar{z}$ are the logical variables in p and q , and the universal quantifier defines the scope of the logical variables (for further details on relational assertions, see e.g. [BvW98, HLL⁺12]). The standard entailment relation is lifted to assertion pairs and to sets of assertion pairs, as follows [DJOS10]:

Definition 4.1 (Entailment). Let (p, q) and (r, t) be assertion pairs and let $\mathcal{U}$ and $\mathcal{V}$ denote the sets $\{(p_i, q_i) \mid 1 \leq i \leq n\}$ and $\{(r_i, t_i) \mid 1 \leq i \leq m\}$. *Entailment* is defined over assertion pairs and sets of assertion pairs by

1. $(p, q) \rightarrow (r, t) \triangleq (\forall \bar{z}_1 . p \Rightarrow q^o) \Rightarrow (\forall \bar{z}_2 . r \Rightarrow t^o)$,
where $\bar{z}_1$ and $\bar{z}_2$ are the logical variables in (p, q) and (r, t) , respectively.
2. $\mathcal{U} \rightarrow (r, t) \triangleq (\bigwedge_{1 \leq i \leq n} (\forall \bar{z}_i . p_i \Rightarrow q_i^o)) \Rightarrow (\forall \bar{z} . r \Rightarrow t^o)$.
3. $\mathcal{U} \rightarrow \mathcal{V} \triangleq \bigwedge_{1 \leq i \leq m} \mathcal{U} \rightarrow (r_i, s_i)$.

The relation $\mathcal{U} \rightarrow (r, t)$ corresponds to classic Hoare style reasoning, proving $\{r\}s\{t\}$ from $\{p_i\}s\{q_i\}$ for all $1 \leq i \leq n$, by means of the consequence and conjunction rules [AdBO09]. When proving entailment, program variables (input and output) are implicitly universally quantified. Furthermore, entailment is reflexive and transitive, and $\mathcal{V} \subseteq \mathcal{U}$ implies $\mathcal{U} \rightarrow \mathcal{V}$.

Behavioral Interfaces and Method Contracts. The (external) behavior of a program can be described in terms of behavioral interfaces which extend the interfaces of a program with method contracts. A *method contract* guarantees that a method has a certain behavior, captured by an assertion pair relating the prestate and the poststate of the method execution. In addition to describing the relation between the values of the method's formal parameters and its returned output value, a method's effect on the state can be described by means of so-called *model variables* (e.g., [HLL⁺12]) in the interface. To avoid representation exposure, a *representation function* is used to relate the actual fields of a class

implementing the interface to the model variables of the interface. Logical variables, which are not model variables, are scoped within an assertion pair so they have the same value in the pre- and postcondition (this scope is formally captured in Definition 4.1 by quantification). In the syntax, an *annotated method header* extends a method header H as defined in Figure 2 with an assertion pair, and an *annotated interface declaration* extends an interface declaration ID with contracts for its exported methods. Formally, annotated method headers AH and annotated interface declarations AID have the syntax

$$\begin{aligned} AH &::= H : (a, a) \\ AID &::= \text{interface } I \text{ extends } \bar{I} \{ \text{model: } \bar{F} \ \bar{AH}; \} \end{aligned}$$

where $\bar{F}$ are the model variables of the interface. To easily distinguish annotations used for behavioral reasoning from the code of the program, italics are used for annotations in the examples, for both model variables and assertions.

Example 4.2. To be concrete in the examples, we consider Boolean assertions a in an assertion language defined by

$$a ::= \text{this} \mid \text{result} \mid \text{null} \mid v \mid z \mid op(\bar{a}).$$

In this assertion language, *this* denotes the current object, *result* the current method’s return value, v a program variable, z a logical variable, and op an operation on data types. In particular, the equality of two assertions a_1 and a_2 is denoted $a_1 == a_2$. We now extend the interfaces $IAcc$ and $IFeeAcc$ of Example 3.1 with contracts as follows:

```
interface IAcc {
  model: int bal, int owner
  void deposit(int x) : (x > 0 ∧ bal == b0, bal = b0 + x);
  void withdraw(int id, int x) :
    (x > 0 ∧ bal == b0, (id == owner ⇒ bal == b0 - x) ∧ (id ≠ owner ⇒ bal == b0));
}

interface IFeeAcc {
  model: int bal, int owner, int fee
  void deposit(int x) : (x > 0 ∧ bal == b0, bal = b0 + x);
  void withdraw(int id, int x) :
    (x > 0 ∧ bal == b0, (id == owner ⇒ bal == b0 - x - fee) ∧ (id ≠ owner ⇒ bal == b0));
}
```

Logical variables are conventionally denoted by subscripts and all formal parameters are read-only. For instance, the initial value of *bal* is captured by the logical variable b_0 in $(bal == b_0, bal == b_0 + x)$, and x is the value of the formal parameter.

Specifying Basic Traits. For a basic trait, a *guarantee* for the behavior of each provided method is given by an assertion pair. Since the guarantee of a provided method m may crucially depend on the behavior of the methods *called* by m , the guarantee of m has an associated set of *requirements* on other methods; e.g., the requirement expressed by $n : (r, t)$ is that the method n must guarantee (r, t) . The guarantee (p, q) of a method together with an associated set $\bar{R}$ of such requirements constitute a *method specification*, denoted $\langle (p, q), \bar{R} \rangle$.

Definition 4.3 (Validity of method specifications). A method specification $\langle (p, q), \bar{R} \rangle$ for a method m with body s is *valid* if the guarantee for s can be derived from the requirements $\bar{R}$ in the program logic PL ; i.e., $\bar{R} \vdash_{PL} \{p\} s \{q\}$.

The validity of a specification $\langle (p, q), \bar{R} \rangle$ for a method m can be mechanically checked by providing a *proof outline* [OG76] for the method body of m . A proof outline typically annotates the method body with information which is difficult to infer but needed for a proof, such as loop invariants and method call annotations. We focus on the latter here, and assume that method calls in a proof outline for the body of m are annotated with the behavioral requirements on auxiliary methods which are needed by m to fulfill the guarantee (p, q) . If external calls are checked with respect to their method contracts as given in behavioral interfaces, it is sufficient for the developer to annotate internal calls. Let $reqs(O)$ denote the set of behavioral requirements to internal method calls in a proof outline O , i.e., each annotated internal call $\{r\} n \{t\}$ in O leads to a requirement $n : (r, t)$ in $reqs(O)$, and let $O \vdash_{PL} s : (p, q)$ denote that O is a *valid proof outline* for the method body s with the guarantee (p, q) ; i.e., $O \vdash_{PL} s : (p, q) \Rightarrow reqs(O) \vdash_{PL} \{p\} s \{q\}$. Generally, it is easier to (mechanically) check $O \vdash_{PL} s : (p, q)$ than to derive a proof of $reqs(O) \vdash_{PL} \{p\} s \{q\}$.

A trait is designed for flexible reuse. It can be difficult to specify the methods in the trait in a way which covers all possible future usages of the trait. In practice, there may be many possible guarantees for the provided methods

of a trait, depending on the context of use. These guarantees can have different associated proof outlines, which give rise to different requirements on the called methods. Thus, a provided method in a trait can have several specifications reflecting different usage contexts. The initial specifications reflect the originally intended usage of the method; further specifications may be added later if new ways of using the trait are discovered. If the initial specification happens to be sufficient for the later actual usage, this is the special case which *coincides with modular specification*. In general, an *annotated method* associates a list of method specifications to a method declaration and an *annotated basic trait* has annotated method declarations for its provided methods. Similarly, an *annotated class* provides a representation function rp_I for each interface I that it implements, which maps model variables f to assertions over the fields of the class. Formally, requirements R , specifications sp , annotated methods AM , annotated basic traits ABT , and annotated classes AC have the syntax

$$\begin{aligned} R &::= n : (a, a) \\ sp &::= \langle (a, a), \bar{R} \rangle \\ AM &::= M \ \overline{sp} \\ ABTE &::= \{ \bar{F}; \bar{H}; \bar{AM} \} \\ ABT &::= \textbf{trait } T_b \textbf{ is } ABTE \\ AC &::= \textbf{class } C \textbf{ implements } \bar{I} \textbf{ by } \{ \overline{rp}_I \ \bar{F}; \} \textbf{ and } CTE \end{aligned}$$

For $ABTE$ we assume a basic trait declaration $\{ \bar{F}; \bar{H}; \bar{M} \}$ such that each AM in $\bar{AM}$ extends a method declaration M in $\bar{M}$ with specifications. The notation introduced for methods can be extended to annotated methods as follows. For a set $\bar{AM}$ of annotated methods, let the function $mtds(\bar{AM})$ return the methods defined in $\bar{AM}$; i.e., $mtds(\bar{AM})$ can be defined by $mtds(\emptyset) \triangleq \emptyset$ and $mtds(M \ \overline{sp} \ \bar{AM}) \triangleq M \ mtds(\bar{AM})$. Then, $m \in \bar{AM}$ is defined by $m \in mtds(\bar{AM})$. Let $\bar{AM}(m)$ return the annotated definition of m in $\bar{AM}$, so $\bar{AM}(m) = AM$ if $m \in \bar{AM}$. The body of an annotated method can be accessed as an unannotated method, $body(M \ \overline{sp}) = body(M)$. Furthermore, $specs(M \ \overline{sp}) \triangleq \overline{sp}$. The specification of methods is illustrated by the following example, where specifications are given in a style similar to JML [BCC⁺05] and Fresco [Wil91].

Example 4.4. Consider the `withdraw` method of trait `TAcc` in Example 3.1. This method may be given the two specifications below. For clarity, a specification $\langle (p, q), \bar{R} \rangle$ is here written **guar:** (p, q) **req:** $\bar{R}$, where the clause **req:** $\bar{R}$ is omitted if there are no requirements. The specifications are labelled $w1$ and $w2$.

```
trait TAcc is {                                // no required fields
  boolean validate(int a); // required method
  void update(int y);      // required method
  void deposit(int x) {update(x);}
  void withdraw(int id, int x) {
    boolean v = validate(id);
    if (v) {update(-x);}
  }
  // w1:
  guar: (bal == b0 ∧ id == owner, bal == b0 - x)
  req:  update: (bal == b0, bal == b0 + y),
        validate: (id == owner, result == true)
  // w2:
  guar: (bal == b0 ∧ id ≠ owner, bal == b0)
  req:  validate: (id ≠ owner, result == false)
}
```

Trivial specifications are here omitted for brevity; e.g., that `bal` is not modified by `validate`. The specifications $w1$ and $w2$ are provided under the assumption that the auxiliary calls to `update` and `validate` manipulate the fields `bal` and `owner`, which are otherwise not required by the trait. However, they are required by the trait `TAux`. `TAux` may be given the following specifications, which do not lead to any requirements:

```
trait TAux is {
  int bal; int owner; // required fields
  void update(int y) {bal = bal + y;}
  guar: (bal == b0, bal == b0 + y)
```

```

boolean validate(int id) {return (id == owner);}
guar: (true,result == (id == owner))
}

```

Specification Entailment. The rule of consequence in Hoare logic [AdBO09] allows us to infer that if $\{p\}s\{q\}$ holds for some statement s and $(p, q) \rightarrow (p', q')$ then $\{p'\}s\{q'\}$. However, the order of entailment is reversed for *substitutability* of assumptions in proof outlines: Given that $\{p\}s_1; s_2; s_3\{q\}$ holds for statements s_1, s_2 , and s_3 by assuming $\{p_1\}s_2\{q_1\}$, let (p_2, q_2) be such that $(p_2, q_2) \rightarrow (p_1, q_1)$. Then $\{p\}s_1; s_2; s_3\{q\}$ must hold by assuming $\{p_2\}s_2\{q_2\}$. This leads to the following observation about substitutability of decorated method calls in proof outlines:

Observation 4.5. Let s be a statement, and p, q, p_1, p_2, q_1, q_2 assertions such that $(p_2, q_2) \rightarrow (p_1, q_1)$. If $A, n() : (p_1, q_1) \vdash \{p\}s\{q\}$ then $A, n() : (p_2, q_2) \vdash \{p\}s\{q\}$.

This observation motivates how the entailment relation is lifted to method specifications to express that a specification $\langle (p_2, q_2), \bar{R}_2 \rangle$ can be *derived* from a proof of a specification $\langle (p_1, q_1), \bar{R}_1 \rangle$.

Definition 4.6 (Specification Entailment). Assume given specifications $\langle (p_1, q_1), \bar{R}_1 \rangle$ and $\langle (p_2, q_2), \bar{R}_2 \rangle$. Specification entailment, denoted $\langle (p_1, q_1), \bar{R}_1 \rangle \rightarrow \langle (p_2, q_2), \bar{R}_2 \rangle$, is defined by

$$\langle (p_1, q_1), \bar{R}_1 \rangle \rightarrow \langle (p_2, q_2), \bar{R}_2 \rangle \triangleq (p_1, q_1) \rightarrow (p_2, q_2) \wedge \bar{R}_2 \rightarrow \bar{R}_1.$$

The following lemma shows that specification entailment preserves the validity of specifications.

Lemma 4.7. Let the specifications $\langle (p_1, q_1), \bar{R}_1 \rangle$ and $\langle (p_2, q_2), \bar{R}_2 \rangle$ be such that $\langle (p_1, q_1), \bar{R}_1 \rangle \rightarrow \langle (p_2, q_2), \bar{R}_2 \rangle$. If $\langle (p_1, q_1), \bar{R}_1 \rangle$ is valid for a method m , then $\langle (p_2, q_2), \bar{R}_2 \rangle$ is also valid for m .

Proof. Follows from the transitivity of specification entailment. $\square$

Assume that the user has supplied $\langle (p_2, q_2), \bar{R}_2 \rangle$ as a specification for some statement s . It follows from Lemma 4.7 that to ensure this specification, we may prove a different specification $\langle (p_1, q_1), \bar{R}_1 \rangle$ with a (possibly) stronger guarantee. Note that the requirement set $\bar{R}_2$ may contain requirements that are superfluous in order to derive the proof outline for (p_1, q_1) .

5. Compositional Verification of Traits

The goal of our verification technique is to reason incrementally about trait expressions while verifying trait-based programs. By *incremental*, we mean that new specifications may be introduced but that old specifications, and consequently old proofs, are never violated. Due to the flexible reuse potential of traits, we do not assume that a fixed specification of a trait, given a priori, covers all potential usages of that trait. This would be a special modular case of our more general incremental approach. Thus, our incremental approach subsumes a modular approach where this is applicable, but it also supports a gradual introduction of specifications in an existing trait library. Instead, traits have a set of possible method specifications for each provided method. This set can be incrementally extended. We devise compositional proof rules that apply to sets of method specifications when traits are composed.

During the verification of trait expressions, a *trait environment* is constructed to keep track of the specifications for the provided methods of trait expressions. The trait environment formalizes a “specification cache” for a trait library, and stores the currently known specifications for methods in a set of traits expressions. The trait environment supports reasoning reuse for the considered trait library in a modular way, but it may also be extended to support incremental reasoning when the known specifications are insufficient. The analysis of a program using the traits of the library interacts with the trait environment to verify the program. For the analysis, we assume that the interfaces of a program have been annotated, so every method in an interface has a method contract as defined in Section 4. This allows us to reason compositionally about calls to methods on interface types. In this section, we examine the construction of the trait environment during the analysis of basic traits and composite traits assembled using the trait composition and alteration operations.

Basic Traits. Consider a basic trait T which provides a method m with the method body t and assume an annotated method declaration $M \vdash \langle (p, q), \bar{R} \rangle$ for m . The correctness of the guarantee (p, q) can be established from the requirement set $\bar{R}$ either by means of automated techniques or by asking the developer for a proof outline O and letting a verification system verify $O \vdash_{PL} t : (p, q)$. For simplicity we assume in the sequel that it is the developer who provides

a proof outline O when necessary. In the proof outline, external calls may be analyzed directly, since they are based on the method contracts in the behavioral interface of the callee. This leaves the requirements to internal calls, which are provided as annotations in the proof outline. However, the behavior of these internal calls depends on how a class is finally assembled from trait expressions. Even if the proof outline is valid, it remains to check that the requirements $\bar{R}$ are correct when the class is assembled. Therefore, the method specification $\langle(p, q), \bar{R}\rangle$ is stored in the trait environment for T . A provided method m may have different guarantees depending on different requirements on its internally called methods. Each of these guarantees is proven using a different proof outline, leading to different specifications for m . Hence, this analysis is repeated for an annotated method AM until the validity of all specifications has been checked.

Trait Modifiers. When a trait alteration expression modifies another trait expression TAE, it gets its own specification stored in the trait environment. This specification is obtained from the specification of TAE and may lead to new proof obligations. *Excluding* a method from a trait does not generate any proof obligations. The trait environment of the resulting trait expression is obtained from the trait environment of the previous trait expression by removing the method specifications of the removed method. *Aliasing* does not generate proof obligations. The trait environment for the resulting trait expression is obtained by copying the method specifications of the aliased method. *Renaming of methods* does not generate proof obligations, but proof obligations for distinct methods may now apply to the same method. The trait environment for the resulting trait expression is obtained by consistently renaming the respective method. *Field renaming* does not generate proof obligations, but some specifications may have to be discarded in order to maintain soundness. Therefore, the trait environment for the trait resulting from the trait alteration $[f \text{ renameTo } f']$ is obtained by distinguishing different cases. If the old field f does not occur in the previous trait expression, the trait environment for the new trait expression is obtained directly from the trait environment of the old trait expression. Otherwise, for each method m , we consider whether the new field f' occurs in the original body of m or not. If f' occurs in m , the specifications of m are dropped, which can be illustrated by the triple $\{b == 5\} a = 2 * b \{a == 10 \wedge b == 5\}$. After $[a \text{ renameTo } b]$ a direct substitution would yield the invalid triple $\{b == 5\} b = 2 * b \{b == 10 \wedge b == 5\}$. If the new variable f' does not occur in the original body of m , specifications containing occurrences of f' are dropped, which can be illustrated by the triple $\{b == 5\} a = 0 \{b == 5\}$. This triple is dropped after the renaming $[a \text{ renameTo } b]$ since the triple $\{b == 5\} b = 0 \{b == 5\}$ is not valid. For the remaining specifications, where f' does not occur in the method body or the specification, new specifications are formed directly by renaming f to f' . This discussion illustrates that field renaming should be used with care in order to avoid unintended name clashes in the program (a more fine-grained analysis of field renaming may be obtained depending on the expressivity of the program logic PL ; e.g., [PBC06]).

Example 5.1. Let the trait $TAux$ be as specified in Example 4.4, and consider the rename operation $TAux[\text{update renameTo basicUpd}]$. The implementation of `validate` is unaltered by the operation, but the `update` method is renamed, and the specification given in $TAux$ applies to the new method:

```
void basicUpd(int y) {bal = bal + y;}
guar: (bal == b0, bal == b0 + y)
```

Symmetric Sum of Traits. For a composed trait expression CTE obtained from two trait expressions by symmetric sum, we maintain the distinction that each method specification has particular assumptions on the required methods. Thus, the trait environment of the composed trait expression is the union of the trait environments of the summand trait expressions. In particular, method specifications are kept in the trait environment even if their requirements cannot become satisfied by the implementations found in other trait expressions of the composition. The reason is that the composed trait expression may be the subject to later trait composition operations. Thus, method specifications that were unsatisfiable in the original composition may again become satisfiable. However, if the composed trait expression is used in a class declaration, the analysis of the class ignores method specifications that are unnecessary in order to verify the interface contract of the class, thereby selecting a set of consistent method specifications from the set of all provided method specifications from the constituents of the composed trait expression.

6. Verifying Trait-Based Programs

For each class, it is necessary to show that every public method exposed through an interface guarantees the contract in that interface. This result should eventually follow from the specification of the methods, as provided by the trait

expression. In the case where the trait specifications contain sufficient guarantees for the public methods, this follows directly in a modular way. Otherwise, a new method specification with the additional guarantees may be added to the trait specifications, and method specifications are collected when a class is assembled from the traits by trait composition. For each added method specification, the respective method must be reinspected to verify the new proof outline associated with the new specification. Such proof outlines may lead to new requirements on internally called methods, which again make it necessary to supply new proof outlines for these methods. This procedure repeats for internal calls until the analysis is complete. Note that all proofs which rely on previously established guarantees of the provided method remain valid. Thus, the presented approach is *incremental*.

Example 6.1. Consider the analysis of the class `CACC` from Example 3.1, which was implemented by

```
class CACC implements IACC
  by {int bal; int owner;} and TACC + TAUX
```

The class must implement the model variables of the interface by providing a representation function (in the terminology of JML [BCC⁺05]); for simplicity in the examples of this paper, the representation function will always just be the identity function and we omit its explicit representation. Thus, in order to implement the interface `IACC`, the method `withdraw` in `TACC` must satisfy the following contract given in Example 4.2: $(x > 0 \wedge bal == b_0, (id == owner \Rightarrow bal == b_0 - x) \wedge (id \neq owner \Rightarrow bal == b_0))$. Since this contract follows by entailment from the guarantees of `w1` and `w2` in Example 4.4, it suffices to ensure that the requirements of `w1` and `w2` are satisfied for the implementations found in `TAUX`, which is straightforward.

Example 6.2. Consider the analysis of the class `CFeeAcc` of Example 3.1, which was implemented by

```
class CFeeAcc implements IFeeAcc
  by {int bal; int owner; int fee;}
  and TACC + TFee + TAUX[update renameTo basicUpd]
```

As for `CACC`, the representation function here is the identity function, and the contract $(x > 0 \wedge bal == b_0, (id == owner \Rightarrow bal == b_0 - x - fee) \wedge (id \neq owner \Rightarrow bal == b_0))$ is imposed by `IFeeAcc` on the method `withdraw`. This contract does not follow from the guarantees of `w1` and `w2`, so a new proof outline is needed. It suffices to extend the specifications of `withdraw` with the following specification, labelled `w3`:

```
void withdraw(int id, int x) {...}
// w3:
guar: (x > 0 ∧ bal == b0 ∧ id == owner, bal == b0 - x - fee)
req:  update: (bal == b0 ∧ y < 0, bal == b0 + y - fee),
      validate: (id == owner, result == true)
```

The interface contract follows by entailment from the guarantees `w2` and `w3`. Now, the requirements of these specifications need to be verified. The only non-trivial requirement is the one to `update`, which can be verified by the following specification in `TFee`:

```
trait TFee is {...}
  void update(int y) {...}
  guar: (y < 0 ∧ bal = b0, bal = b0 + y - fee)
  req:  basicUpd: (bal = b0, bal = b0 + y)
```

The requirement of this specification follows from the guarantee of `basicUpd` as given in Example 5.1.

Example 6.3. Assume that the trait `TACC` is specified as in Example 4.4. We define trait `TMini` by

```
trait TMini is {
  int bal; // required field
  void update(int y) {bal=bal+y;}
  void validate(int id) {return true;}
}
```

Assume that class `CMini` is defined by `TACC+TMini` where `bal` is the only declared field. The original specifications `w1` and `w2` contain the field `owner`, reflecting that `TACC` was originally developed together with `TAUX`, but this field is not present in the current composition. To analyze method `withdraw`, we may provide the following specification:

```

void withdraw(int id, int x) {...}
// w4:
guar: (bal == b0, bal == b0 - x)
req:  update : (bal == b0, bal == b0 - x)
       validate : (true, result == true)

```

This specification expresses that any client can make a withdrawal from a CMiniAcc.

7. PST(PL): A Proof System for Verifying Trait-Based Programs

The verification process for trait-based programs outlined in the previous sections is now formalized as a calculus PST(PL) which is parametric in the underlying program logic PL . The calculus PST(PL) relies on a sound program logic PL and defines inference rules for analyzing trait expressions and classes, given in Sections 7.2 and 7.3 below. The proof system considers annotated basic traits ABT. For simplicity, we assume that unannotated method declarations in AM have a default guarantee $(true, true)$, with corresponding requirements $(true, true)$ for internal calls.

A program is analyzed in PST(PL) as a sequence of analysis operations concerning the trait and class definitions of the program. These operations manipulate a *proof environment* which consists of a trait environment and a class environment. For each analyzed trait alteration expression, the trait environment is extended with the trait definition and with the specifications for the methods provided by that trait expression. Thus, if the analysis of a trait expression TAE is initiated in some trait environment $\mathcal{T}$, the successful analysis of TAE will lead to a trait environment $\mathcal{T}'$, such that $\mathcal{T}'$ extends $\mathcal{T}$. In this case we say that $\mathcal{T}'$ is the trait environment *resulting* from the analysis. When analyzing classes, the class environment is extended similarly. Traits and classes are analyzed in the context of the trait and class environments which have been obtained from the analysis of previous traits and classes. In this way, the rules of PST(PL) explain how the analysis of language artefacts constructs a sequence of proof environments, starting with empty trait and class environments. This process formalizes the accumulation of specifications for a collection of traits during the analysis of programs using those traits.

7.1. Judgments in PST(PL)

Judgments in PST(PL) are of the form $\mathcal{C}, \mathcal{T} \vdash \mathcal{P}$, where $\mathcal{C}$ is a *class environment* for class analysis, $\mathcal{T}$ is a *trait environment* for trait analysis, and $\mathcal{P}$ is a sequence of *analysis operations*. We define the syntax for judgments in PST(PL), including the analysis operations and the proof environment which is manipulated by these operations.

7.1.1. Proof Environments

Class environments are used to accumulate knowledge about classes during the analysis. Class environments, which represent classes by a unique name and a tuple $\langle \bar{I}, \text{CTE}, \bar{F} \rangle$ of type *Class*, are defined as follows:

Definition 7.1 (Class environments). A class environment $\mathcal{C}$ consists of two mappings $D_{\mathcal{C}}$ and $S_{\mathcal{C}}$, where $D_{\mathcal{C}} : \text{Cid} \mapsto \text{Class}$, and $S_{\mathcal{C}} : \text{Cid} \times \text{Mid} \mapsto \text{Set}[\text{Spec}]$.

Here, $D_{\mathcal{C}}$ contains the definitions of verified classes and $S_{\mathcal{C}}$ contains their verified specifications. The main purpose of the class environment is to record the method specifications used to establish the contracts of the implemented interfaces. The function $\text{impl} : \text{Class} \rightarrow \text{CTE}$ returns the trait expression of a class. Update functions for the class environment $\mathcal{C}$ are defined as follows, where $\bar{\text{sp}}$ is a set of method specifications:

$$\begin{aligned}
 \mathcal{C} \oplus (\text{C}, \bar{I}, \text{CTE}, \bar{r}\bar{p}_I, \bar{F}) &\triangleq \langle D_{\mathcal{C}}[\text{C} \mapsto \langle \bar{I}, \text{CTE}, \bar{r}\bar{p}_I, \bar{F} \rangle], S_{\mathcal{C}} \rangle \\
 \mathcal{C} \oplus (\text{C}, \text{m}, \bar{\text{sp}}) &\triangleq \langle D_{\mathcal{C}}, S_{\mathcal{C}}[(\text{C}, \text{m}) \mapsto S_{\mathcal{C}}(\text{C}, \text{m}) \cup \bar{\text{sp}}] \rangle
 \end{aligned}$$

Trait environments are used to accumulate knowledge about trait expressions during analysis, and are defined as follows:

Definition 7.2 (Trait environments). A trait environment $\mathcal{T}$ is a mapping from trait alteration expressions to annotated basic traits: $\mathcal{T} : \text{TAE} \rightarrow \text{ABTE}$.

Recall that trait alteration expressions TAE are defined by $\text{Tb } \bar{\text{a}}\bar{\text{o}}$; i.e., they consist of a basic trait name followed by a possibly empty sequence of alteration operations. The mapping $\mathcal{T}$ takes such a trait alteration expression and returns

$$\begin{aligned}
\mathcal{P} &::= \text{trait } T \text{ is } \{\bar{F}; \bar{H}; \bar{AM}\} \mid \text{trait } T \text{ is } CTE \mid \text{verify}(T, \bar{AM}) \mid \text{extend}(CTE) \\
&\quad \mid \text{class } C \text{ implements } \bar{T} \text{ by } \{\bar{r}\bar{p}_I \bar{F};\} \text{ and } CTE \mid \langle C : \mathcal{O} \rangle \mid \mathcal{P} \cdot \mathcal{P} \\
\mathcal{O} &::= \varepsilon \mid \text{discharge}(\bar{R}) \mid \text{analyze}(TAE, m : (a, a)) \mid \mathcal{O} \cdot \mathcal{O}
\end{aligned}$$

Fig. 5. Syntax of the analysis operations.

the annotated implementation of the trait as a basic trait expression. In this manner, the trait environment is used to capture the provided methods of the different traits and to associate a number of specifications with each method.

Auxiliary functions on trait environments. For a trait alteration expression TAE, we let $TAE \in \mathcal{T}$ denote that TAE is in the domain of $\mathcal{T}$, so $\mathcal{T}(TAE)$ is defined. For $\mathcal{T}(TAE) = \{\bar{F}; \bar{H}; \bar{AM}\}$ we lift functions over annotated method sets such that $m \in \mathcal{T}(TAE) \triangleq m \in \bar{AM}$, $\mathcal{T}(TAE)(m) \triangleq \bar{AM}(m)$, and $mtds(\mathcal{T}(TAE)) \triangleq mtds(\bar{AM})$. For a composed trait expression CTE defined by $TAE_1 + \dots + TAE_n$ where each $TAE_i \in \mathcal{T}$, we define $mtds(CTE)$ straightforwardly by the union $\bigcup_{1 \leq i \leq n} mtds(\mathcal{T}(TAE_i))$. By type safety, we may assume that name conflicts do not occur in the symmetric sum; i.e., each method in a sum of trait expressions is defined in exactly one summand TAE_i . Let the function $addSpec$ be defined such that $addSpec(\bar{AM}, m, sp)$ returns $\bar{AM}$ where $\bar{AM}(m)$ is extended by sp , and let $addSpec(\{\bar{F}; \bar{H}; \bar{AM}\}, m, sp) \triangleq \{\bar{F}; \bar{H}; addSpec(\bar{AM}, m, sp)\}$.

The *update function* $\oplus$ is used to extend environments with new definitions and specifications during the analysis. Let T be the name of a basic trait defined by $\{\bar{F}; \bar{H}; \bar{AM}\}$, m a method name, and $\langle (p, q), \bar{R} \rangle$ a specification. A trait environment $\mathcal{T}$ is extended with new basic traits and specifications as follows:

$$\begin{aligned}
\mathcal{T} \oplus (T, \{\bar{F}; \bar{H}; \bar{AM}\}) &\triangleq \mathcal{T}[T \mapsto \{\bar{F}; \bar{H}; \bar{AM}\}] \\
\mathcal{T} \oplus (T, m, sp) &\triangleq \mathcal{T}[T \mapsto addSpec(\mathcal{T}(T), m, sp)]
\end{aligned}$$

7.1.2. Analysis operations

The analysis assumes that basic traits are annotated as explained in Section 4. The syntax for the analysis operations $\mathcal{P}$ is given in Figure 5. Analysis operations include user-given trait and class definitions, in addition to $verify(T, \bar{AM})$, $extend(CTE)$, and $\langle C : \mathcal{O} \rangle$. The operation $verify(T, \bar{AM})$ applies to a trait T , where $\bar{AM}$ is a set of methods annotated with specifications to be verified for the trait. For a composed trait expression CTE, the operation $extend(CTE)$ is used to construct an entry in the trait environment for each trait alteration expression in CTE. Finally, operation $\langle C : \mathcal{O} \rangle$ encapsulates the operations $\mathcal{O}$, ensuring that they are analyzed within the context of a specific class C . Each element in $\mathcal{O}$ is either a $discharge(\bar{R})$ operation, which indicates that the requirement set $\bar{R}$ must be verified for C , or an $analyze(TAE, n : (r, s))$ operation, which indicates that the requirement $n : (r, s)$ must be further analyzed in the context of class C , where the definition of n stems from the trait expression TAE.

7.2. Trait Analysis in PST(PL)

The inference rules for trait analysis are given in Figure 6. Traits are analyzed compositionally in the context of the trait environment $\mathcal{T}$. A trait is either an annotated basic trait of the form **trait** T_b **is** $\{\bar{F}; \bar{H}; \bar{AM}\}$ or a composed trait of the form **trait** T_c **is** CTE. For an annotated basic trait $\{\bar{F}; \bar{H}; \bar{AM}\}$ we assume by type safety that $\bar{F}$ and $\bar{H}$ contain all fields and method signatures used in the provided methods. Furthermore, we assume for simplicity that $\bar{F}$ and $\bar{H}$ are *minimal* in the sense that they do not contain field and method signatures that are not used in the provided methods.

An annotated basic trait **trait** T_b **is** $\{\bar{F}; \bar{H}; \bar{AM}\}$ is analyzed by the rule **BASICTRAIT**, which extends $\mathcal{T}$ with the definition of the trait and generates a *verify* operation for analyzing the user-given specifications for the provided methods of the trait. The *verify* operation on a set $\bar{AM}$ of annotated methods is decomposed by the rules **DECOMP1** and **DECOMP2**, which leads to a $verify(T_b, m \text{ } sp)$ operation for each specification sp of $\text{Min } \bar{AM}$. The predicates $\text{nonempty}(\bar{AM}_i)$ and $\text{nonempty}(\bar{sp}_i)$ (for $i \in \{1, 2\}$) express that decomposition rules only apply to nonempty sequences. The rule **ADAPTATION** builds specification entailment into the proof system. This rule replaces the user-provided specification by a stronger specification. For a method $\bar{m}(\bar{T} \bar{x})\{t\}$ with the specification $\langle (p, q), \bar{R} \rangle$, the rule **VERIFY** requires that the user provides a proof outline O for the method body t such that $O \vdash_{PL} t : (p, q)$ (for example by the user, as discussed in Section 5). In this case $\langle (p, q), reqs(O) \rangle$ is a valid method specification for m , and the annotated method $\mathcal{T}(T_b)(m)$ in the trait environment is extended with this specification.

Upon the successful analysis of T_b , each specification is recorded by the trait environment $\mathcal{T}$. Note that when

$$\begin{array}{c}
\text{(BASICTRAIT)} \\
\frac{Tb \notin \mathcal{T} \quad \mathcal{C}, \mathcal{T} \oplus (Tb, \{\bar{F}; \bar{H}; \bar{A}M\}) \vdash \text{verify}(Tb, \bar{A}M) \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \mathbf{trait} Tb \mathbf{is} \{\bar{F}; \bar{H}; \bar{A}M\} \cdot \mathcal{P}} \\
\\
\begin{array}{cc}
\text{(DECOMP1)} & \text{(DECOMP2)} \\
\frac{\text{nonempty}(\bar{A}M_1) \quad \text{nonempty}(\bar{A}M_2)}{\mathcal{C}, \mathcal{T} \vdash \text{verify}(T, \bar{A}M_1) \cdot \text{verify}(T, \bar{A}M_2) \cdot \mathcal{P}} & \frac{\text{nonempty}(\bar{s}p_1) \quad \text{nonempty}(\bar{s}p_2)}{\mathcal{C}, \mathcal{T} \vdash \text{verify}(T, M \bar{s}p_1) \cdot \text{verify}(T, M \bar{s}p_2) \cdot \mathcal{P}} \\
\mathcal{C}, \mathcal{T} \vdash \text{verify}(T, \bar{A}M_1 \bar{A}M_2) \cdot \mathcal{P} & \mathcal{C}, \mathcal{T} \vdash \text{verify}(T, M \bar{s}p_1 \bar{s}p_2) \cdot \mathcal{P}
\end{array} \\
\\
\text{(ADAPTATION)} \\
\frac{\langle (p', q'), \bar{R}' \rangle \rightarrow \langle (p, q), \bar{R} \rangle \quad \mathcal{C}, \mathcal{T} \vdash \text{verify}(T, M \langle (p', q'), \bar{R}' \rangle) \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \text{verify}(T, M \langle (p, q), \bar{R} \rangle) \cdot \mathcal{P}} \\
\\
\text{(VERIFY)} \\
\frac{O \vdash_{PL} t : (p, q) \quad \text{reqs}(O) = \bar{R} \quad \mathcal{C}, \mathcal{T} \oplus (T, m, \langle (p, q), \bar{R} \rangle) \vdash \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \text{verify}(T, I m(\bar{I} \bar{x})\{t\} \langle (p, q), \bar{R} \rangle) \cdot \mathcal{P}} \\
\\
\begin{array}{cc}
\text{(COMPTRAIT)} & \text{(DECOMP3)} \\
\frac{\mathcal{C}, \mathcal{T} \vdash \text{extend}(CTE) \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \mathbf{trait} Tc \mathbf{is} CTE \cdot \mathcal{P}} & \frac{\mathcal{C}, \mathcal{T} \vdash \text{extend}(TAE) \cdot \text{extend}(CTE) \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \text{extend}(TAE + CTE) \cdot \mathcal{P}}
\end{array} \\
\\
\begin{array}{cc}
\text{(LOOKUP)} & \text{(EXTEND)} \\
\frac{TAE \in \mathcal{T} \quad \mathcal{C}, \mathcal{T} \vdash \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \text{extend}(TAE) \cdot \mathcal{P}} & \frac{TAE \in \mathcal{T} \quad TAE \text{ ao} \notin \mathcal{T} \quad \mathcal{C}, \mathcal{T} \oplus (TAE, \text{ao}) \vdash \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \text{extend}(TAE \text{ ao}) \cdot \mathcal{P}}
\end{array} \\
\\
\text{(EXTENDREC)} \\
\frac{TAE \notin \mathcal{T} \quad \mathcal{C}, \mathcal{T} \vdash \text{extend}(TAE) \cdot \text{extend}(TAE \text{ ao}) \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \text{extend}(TAE \text{ ao}) \cdot \mathcal{P}}
\end{array}$$

Fig. 6. The inference rules for trait analysis.

the trait is defined it is not known to which actual implementation an internal call will be bound, since the method binding depends on how the traits are used to form classes. Consequently, the imposed requirements are not verified with regard to any implementation during trait analysis. Instead these requirements are verified as needed when a specification is actually used during the analysis of a class.

Composed traits **trait** Tc **is** CTE are analyzed by the rule COMPTRAIT. Here CTE is of the form $TAE_1 + \dots + TAE_n$ for $n \geq 1$, where each TAE_i is of the form $Tb_i \bar{a}o_i$. For simplicity, we here assume that all basic traits $Tb_1, \dots, Tb_n$ have been analyzed before the analysis of composite trait expressions. By the rule COMPTRAIT, an operation *extend* is generated in order to extend the trait environment for each trait alteration expression TAE_i in the composition. The composed trait expression is decomposed to a sequence of *extend* operations by the rule DECOMP3. Rule LOOKUP applies to TAE_i if TAE_i is already defined in the trait environment. Especially, this rule applies to all basic trait names used in the composition. Otherwise, the analysis depends on the structure of TAE_i , and the *extend* operation is analyzed by either EXTEND or EXTENDREC.

Rule EXTEND is the main rule for extending the trait environment for each TAE_i of the form $TAE \text{ ao}$. The premise of the rule ensures that the prefix TAE is defined in the trait environment. The trait environment is extended for $TAE \text{ ao}$ by the update $\mathcal{T} \oplus (TAE, \text{ao})$, which modifies a copy of the annotated basic trait bound to TAE depending on the actual alteration operation ao . The definition of this update function is given in Appendix A. For an expression $TAE \text{ ao}$ where TAE has not already been analyzed (i.e., TAE is not in $\mathcal{T}$), the rule EXTENDREC recursively ensures that the trait environment is extended for TAE before it is extended for $TAE \text{ ao}$. After extending the trait environment for TAE , the rule EXTEND may be applied to *extend*($TAE \text{ ao}$). Since TAE starts with a basic trait name Tb , the recursion over the structure of $TAE \text{ ao}$ is guaranteed to terminate. In this manner, properties for each TAE_i in the composition are remembered in the trait environment. By the successful analysis of Tc , the mapping $\mathcal{T}$ binds TAE_i to an annotated basic trait definition containing the methods provided by TAE_i with verified specifications. These annotated methods are derived by manipulating those of Tb_i according to the modifiers $\bar{a}o_i$.

A trait alteration expression TAE is introduced into $\mathcal{T}$ by rule BASICTRAIT or rule EXTEND. After the initial analysis by one of these rules, the method definitions in TAE are not manipulated, but specifications may be added.

$$\begin{array}{c}
\text{(CLASS)} \\
\frac{\mathcal{C} \oplus (C, \bar{I}, \text{CTE}, \overline{rp_I}, \bar{F}), \mathcal{T} \vdash \text{extend}(\text{CTE}) \cdot \langle C : \text{discharge}(\text{contracts}(\overline{rp_I}, \bar{I})) \rangle \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \text{class } C \text{ implements } \bar{I} \text{ by } \{ \overline{rp_I} \bar{F}; \} \text{ and CTE} \cdot \mathcal{P}} \\
\\
\text{(DECOMP4)} \\
\frac{\mathcal{C}, \mathcal{T} \vdash \langle C : \text{discharge}(\bar{R}_1) \cdot \text{discharge}(\bar{R}_2) \cdot \mathcal{O} \rangle \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \langle C : \text{discharge}(\bar{R}_1 \cup \bar{R}_2) \cdot \mathcal{O} \rangle \cdot \mathcal{P}} \\
\\
\text{(OPENANALYSIS)} \\
\frac{\text{TAE} \in \text{impl}(D_{\mathcal{C}}(C)) \quad m \in \mathcal{T}(\text{TAE}) \quad \mathcal{C}, \mathcal{T} \vdash \langle C : \text{analyze}(\text{TAE}, m : (p, q)) \cdot \mathcal{O} \rangle \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \langle C : \text{discharge}(m : (p, q)) \cdot \mathcal{O} \rangle \cdot \mathcal{P}} \\
\\
\text{(CLOSEANALYSIS)} \\
\frac{\text{guar}(S_{\mathcal{C}}(C, m)) \rightarrow (p, q) \quad \mathcal{C}, \mathcal{T} \vdash \langle C : \mathcal{O} \rangle \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \langle C : \text{discharge}(m : (p, q)) \cdot \mathcal{O} \rangle \cdot \mathcal{P}} \\
\\
\text{(ANALYZE)} \\
\frac{\overline{sp} \subseteq \text{specs}(\mathcal{T}(\text{TAE})(m)) \quad \text{guar}(\overline{sp}) \rightarrow (p, q) \quad \mathcal{C} \oplus (C, m, \overline{sp}), \mathcal{T} \vdash \langle C : \text{discharge}(\text{req}(\overline{sp})) \cdot \mathcal{O} \rangle \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \langle C : \text{analyze}(\text{TAE}, m : (p, q)) \cdot \mathcal{O} \rangle \cdot \mathcal{P}} \\
\\
\text{(INCREMENT)} \\
\frac{\mathcal{C}, \mathcal{T} \vdash \text{verify}(\text{TAE}, \mathcal{T}(\text{TAE})(m)((p, q), \bar{R})) \cdot \langle C : \text{analyze}(\text{TAE}, m : (p, q)) \cdot \mathcal{O} \rangle \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \langle C : \text{analyze}(\text{TAE}, m : (p, q)) \cdot \mathcal{O} \rangle \cdot \mathcal{P}} \\
\\
\text{(EMPDISCHARGE)} \qquad \text{(EMPClass)} \\
\frac{\mathcal{C}, \mathcal{T} \vdash \langle C : \mathcal{O} \rangle \cdot \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \langle C : \text{discharge}(\emptyset) \cdot \mathcal{O} \rangle \cdot \mathcal{P}} \qquad \frac{\mathcal{C}, \mathcal{T} \vdash \mathcal{P}}{\mathcal{C}, \mathcal{T} \vdash \langle C : \varepsilon \rangle \cdot \mathcal{P}}
\end{array}$$

Fig. 7. The inference rules for class analysis. Here, $\text{TAE} \in \text{CTE}$ denotes that TAE is a syntactical part of CTE , i.e., if CTE is $\text{TAE}_1 + \dots + \text{TAE}_n$, then $\text{TAE} = \text{TAE}_i$ for some i ($1 \leq i \leq n$).

The inference system for trait analysis ensures that there exists a valid proof outline for each specification recorded by the trait environment. (This is formalized by Lemma B.2 in Appendix B.) Especially, the analysis ensures that each user given specification in a basic trait is valid.

7.3. Class Analysis in PST(PL)

We now consider the analysis of a class declaration **class** C **implements** $\bar{I}$ **by** $\{ \overline{rp_I} \bar{F}; \}$ **and** CTE . As for composed traits, the expression CTE is of the form $\text{TAE}_1 + \dots + \text{TAE}_n$ and by type safety we have that each provided method m is defined in exactly one of the summand trait expressions TAE_i . In addition to the trait environment $\mathcal{T}$, the class analysis extends a class environment $\mathcal{C}$ which contains the definitions and specifications of classes. The analysis of class C is driven by the contracts of its interfaces $\bar{I}$. Upon the successful analysis of C , each contract of a provided method m in the interfaces of the class follows by *entailment* from the guarantees of $S_{\mathcal{C}}(C, m)$. If m is provided by TAE_i , the interface contracts are ensured by reusing already verified specifications contained in $\mathcal{T}(\text{TAE}_i)(m)$, and by extending this set if needed. Thus, $S_{\mathcal{C}}(C, m)$ contains the subset of $\mathcal{T}(\text{TAE}_i)(m)$ that is actually used to verify the current class. In addition, the requirements imposed by the used specifications are analyzed with regard to the implementation to which they bind in C . For a method specification $\text{sp} = \langle (p, q), \bar{R} \rangle$, we define the functions $\text{guar} : \text{Spec} \rightarrow \text{Guar}$ and $\text{req} : \text{Spec} \rightarrow \text{Set}[\text{Req}]$ where $\text{guar}(\text{sp}) \triangleq (p, q)$ and $\text{req}(\text{sp}) \triangleq \bar{R}$. These functions are straightforwardly lifted to sets of method specifications, returning sets of guarantees and requirements, respectively. Thus, if $\text{sp} \in S_{\mathcal{C}}(C, m)$, then each requirement $n : (r, s) \in \text{req}(\text{sp})$ follows by entailment from the guarantees of $S_{\mathcal{C}}(C, n)$.

Let $\text{contracts}(\overline{rp_I}, \bar{I})$ denote the set of method contracts of $\bar{I}$ obtained by applying the respective representation functions $\overline{rp_I}$ to the model variables in the assertions of each interface, which are of the form $m : (p, q)$. The analysis of the class C is initiated by the rule CLASS. By this rule, the class environment $\mathcal{C}$ is extended by the definition of C , and the analysis operation $\langle C : \text{discharge}(\text{contracts}(\overline{rp_I}, \bar{I})) \rangle$ is generated. This operation reflects that the analysis of C

is driven by the contracts of the implemented interfaces, in which the model fields of the contracts have been adapted to the class by application of its representation function for each interface, and that these contracts are analyzed in the context of C . The set of method contracts is decomposed by the rule **DECOMP4**, and each contract $m : (p, q)$ is analyzed either by rule **OPENANALYSIS** or **CLOSEANALYSIS**. The rule **CLOSEANALYSIS** applies to a *discharge* operation if the requirement follows from earlier verified specifications for the current class. Otherwise, the requirement is opened for analysis at the trait level by rule **OPENANALYSIS**. This rule selects the conjunct TAE in CTE where m is defined, leading to an operation $analyze(TAE, m : (p, q))$. This operation is analyzed by rule **ANALYZE**.

Rule **ANALYZE** captures *modular reasoning* in the inference system, reusing previous analysis from the trait environment. This rule applies when the contract can be derived from previously verified method specifications for m in TAE . In that case, the analysis continues with an *discharge* operation for the requirements of these specifications, which are analyzed in the same manner as the original interface contract. Consequently, each requirement is either discharged by **CLOSEANALYSIS**, or analyzed using **OPENANALYSIS** with respect to the actual implementation of the called method in the current trait composition CTE .

The rule **INCREMENT** captures *incremental reasoning* in the inference system and allows the user to extend the verified specifications of a trait in the trait environment, incrementally building knowledge about the traits. It needs to be applied if the requirement of an *analyze* operation does not follow from previously verified specifications of the current trait. The user must suggest the additional specification $\langle (p, q), \bar{R} \rangle$ for the trait, which is analyzed by a *verify* operation for the current trait. This *verify* operation is resolved by the inference rules for trait analysis and the trait environment is extended before the class analysis proceeds. To cover the case of self-recursion, the method specification itself is assumed when analyzing the requirements, similar to rule **BASICTRAIT** above. This captures the standard approach to reasoning about recursive method calls [Hoa71]. An empty set of requirements or contracts is discarded by the rule **EMPDISCHARGE**, and the successful analysis of class C is completed by the rule **EMPCCLASS**.

7.4. Soundness of PST(PL)

When reasoning about a set of mutually recursive methods, the guarantees in the specifications of all methods are assumed to hold in order to verify the body of each method (e.g., [AdBO09]). We now extend this approach to define the *consistency* of a set of proof outlines for methods in a flattened class with given interfaces. The flattened version of a class that is defined by **class** C **implements** $\bar{I}$ **by** $\{\bar{r}\bar{p}_I, \bar{F};\}$ **and** CTE is given by **class** C **implements** $\bar{I} \{ \bar{r}\bar{p}_I, \bar{F}; \bar{M} \}$ as defined in Section 3.1 (where the annotation with representation functions is unchanged by the flattening).

Definition 7.3 (Consistency). Consider a flattened class **class** C **implements** $\bar{I} \{ \bar{r}\bar{p}_I, \bar{F}; \bar{M} \}$. For each method $m \in \bar{M}$ with method body t , let S_m be a set of method specifications such that for each $\langle (p, q), \bar{R} \rangle \in S_m$, there exists a proof outline O where $O \vdash_{PL} t : (p, q)$ and $\bar{R} = reqs(O)$. The specifications $S_{\bar{M}}$ are *consistent* iff, for all $m \in \bar{M}$:

1. $\forall (m : (r, s)) \in contracts(\bar{r}\bar{p}_I, \bar{I}) . guar(S_m) \rightarrow (r, s)$
2. $\forall \langle (p, q), \bar{R} \rangle \in S_m . \forall (n : (r, s)) \in \bar{R} . guar(S_n) \rightarrow (r, s)$

The first condition expresses that the interface contracts are satisfied, whereas the second condition expresses that the requirements of all internal calls are satisfied. Previous work [DJOS10] defines a sound calculus for analyzing single inheritance class hierarchies. Given a consistent set of specifications, the analysis of a flattened class succeeds in this calculus. In order to ensure soundness of PST(PL), it thereby suffices to prove that the successful analysis of some class C leads to a consistent set of specifications for the flattened version of C .

Theorem 7.4 (Soundness of Trait Verification). For a given class **class** C **implements** $\bar{I}$ **by** $\{\bar{F};\}$ **and** CTE , if the successful analysis of C in PST(PL) results in a class environment $\mathcal{C}$, then the set of method specifications for C in $\mathcal{C}$ are consistent with the flattened version of C .

The proof of this theorem is given in Appendix B.

8. Example: A Trait-Based Implementation of Bank Accounts

This section illustrates how traits can be used to construct a range of classes implementing behavioral interfaces which specify different bank accounts, and how these classes can be verified by means of the proposed incremental proof system. We start with the trait definitions given in Figure 8, where each trait contains a set of annotated methods. To

```

trait TBasicAccount is {
  int bal;           // required field
  void update(nat x); // required method
  void deposit(nat x) {bal=bal+x;} <((bal == b0, bal == b0 + x), 0)>
  void withdraw(nat x) {update(x);}
}

trait TBasicUpd is {
  int bal;           // required field, no required methods
  void update(nat x) {bal=bal-x;} <((bal == b0, bal == b0 - x), 0)>
}

trait TFeeUpd is {
  nat fee;           // required field
  void bUpdate(nat x); // required method
  void update(nat x) {bUpdate(x+fee)}
}

trait TPosUpd is {
  int bal;           // required field
  void bUpdate(nat x); // required method
  void update(nat x) {if (bal >= x) {bUpdate(x);} }
}

```

Fig. 8. Basic trait definitions for implementing bank accounts.

simplify the presentation, we assume that the language provides the type **nat** of positive integers. Default specifications, with guarantees (*true, true*), are omitted. Let us assume that only a few specifications are initially provided in the traits. The incremental aspect of PST(PL) is used to provide additional specifications to the traits when these are required by the class analysis.

The trait **TBasicAccount** implements basic functionality for bank accounts by the two methods **deposit** and **withdraw**; the method **deposit** increases the balance **bal** of the account by the parameter value and **withdraw** decreases the balance of the account by calling an auxiliary method **update**. The exact definition of **update** depends on the specific account, so **update** is a required method in this trait. Additionally, **deposit** assumes access to a field **bal**, so **bal** is also required by the trait. To implement an account, the trait **TBasicAccount** must be composed with another trait where **update** is defined, in a class where **bal** is defined.

The trait **TBasicUpd** provides a basic implementation of the **update** method which simply reduces the balance by the argument value. In contrast, the trait **TFeeUpd** charges a fee for each update, and **update** is defined by calling an auxiliary, required method **bUpdate** where the argument is increased by **fee**. In a similar way, the trait **TPosUpd** calls an auxiliary, required method **bUpdate**. However, in this case the required method in **TPosUpd** is only called if the current balance is greater or equal to the parameter value.

Bank accounts with different behavioral properties may be assembled, depending on how the traits of Figure 8 are combined. Different trait combinations lead to different behaviors for the **withdraw** method. These bank accounts are first specified in terms of behavioral interfaces and then implemented in classes using different trait combinations. We define an interface hierarchy such that each type of account extends a superinterface **IAccount** which declares the methods **deposit** and **withdraw**, but where **withdraw** has no behavioral requirement. This interface further specifies the effect of making a deposit as increasing the balance of the bank account by the deposited amount. The subinterfaces vary in the properties they specify for **withdraw**. In the interface **IBasicAccount**, **withdraw** decreases the balance of the bank account by *exactly* the withdrawn amount; in **IFeeAccount** the balance is decreased by a *fee* on withdrawals; in **IPosAccount** the account cannot be overdrawn; in **IPosFeeAccount** a fee is charged on withdrawal and the withdrawn amount cannot be larger than the balance on the account; and in **IFeePosAccount**, a fee is charged on withdrawals and the balance is guaranteed not to be overdrawn. The definitions of these interfaces are given in Figure 9; observe that the method **deposit** is only specified in **IAccount** whereas **withdraw** is given different specifications in the different subinterfaces of **IAccount**.

The interfaces are implemented by classes, shown in Figure 10, which combine **TBasicAccount** with other

```

interface IAccount {
  model: int bal
  void deposit(nat x) : (bal == b0, bal == b0 + x)
  void withdraw(nat x) // no behavioral requirement
}

interface IBasicAccount extends IAccount {
  model: int bal
  void withdraw(nat x) : (bal == b0, bal == b0 - x)
}

interface IFeeAccount extends IAccount {
  model: int bal, int fee
  void withdraw(nat x) : (bal == b0, bal == b0 - x - fee)
}

interface IPosAccount extends IAccount {
  model: int bal
  void withdraw(nat x) :
    (bal == b0, (b0 ≥ x ⇒ bal == b0 - x) ∧ (b0 < x ⇒ bal == b0),
    (bal ≥ 0, bal ≥ 0))
}

interface IPosFeeAccount extends IAccount {
  model: int bal, int fee
  void withdraw(nat x) :
    (bal == b0, (b0 ≥ x ⇒ bal == b0 - x - fee) ∧ (b0 < x ⇒ bal == b0))
    (bal ≥ 0, bal ≥ -fee)
}

interface IFeePosAccount extends IAccount {
  model: int bal, int fee
  void withdraw(nat x) :
    (bal == b0, (b0 ≥ (x + fee) ⇒ bal == b0 - x - fee) ∧ (b0 < (x + fee) ⇒ bal == b0))
    (bal ≥ 0, bal ≥ 0)
}

```

Fig. 9. Interface declarations for different bank accounts.

traits in different ways. The update method is renamed in some of the traits such that the call to update in withdraw may lead to a chain of calls to renamed update methods. The class CBasicAccount is defined in terms of the basic update method. The class CFeeAccount combines the basic update method with the one from TFeeUpd, so the balance is reduced by an additional fee for each withdrawal. For comparison, the flattened version of this class is shown in Figure 11. The class CPosAccount is defined in a similar way, but the update is not performed if it would lead to a negative balance. The classes CPosFeeAccount and CFeePosAccount, combine all the traits TFeeUpd, TPosUpd, and TBasicUpd. Here, the names reflect the binding order of a call to update. In CPosFeeAccount, a call to update will bind to the implementation found in TPosUpd and a recursive call to the implementation in TFeeUpd. Thus, in CPosFeeAccount the balance must be positive before the fee is added to the argument. This means that $-fee$ is the lower limit of the balance, as specified by the interface IPosFeeAccount. In contrast, the traits of CFeePosAccount are ordered such that the fee is added to the argument before the positive balance test. In this way, CFeePosAccount ensures that the balance is non-negative after a call to update. The verification of these classes with respect to their interfaces is discussed in Appendix C.

Figure 12 summarizes the environments $\mathcal{C}, \mathcal{I}$ resulting from the analysis of the different classes, focussing on the method withdraw. We assume that classes are analyzed in the top-down order in which they appear in the table, after the analysis of the basic trait definitions in Figure 8. Thus, the analysis is assumed to start in some initial environments

```

class CBasicAccount implements IBasicAccount
  by {int bal;} and TBasicAccount + TBasicUpd

class CFeeAccount implements IFeeAccount
  by {int bal, nat fee;}
  and TBasicAccount + TFeeUpd + TBasicUpd[update renameTo bUpdate]

class CPosAccount implements IPosAccount
  by {int bal;}
  and TBasicAccount + TPosUpd + TBasicUpd[update renameTo bUpdate]

class CPosFeeAccount implements IPosFeeAccount
  by {int bal, nat fee;}
  and TBasicAccount + TPosUpd + TBasicUpd[update renameTo bbUpdate]
  + TFeeUpd[bUpdate renameTo bbUpdate][update renameTo bUpdate]

class CFeePosAccount implements IFeePosAccount
  by {int bal, nat fee;}
  and TBasicAccount + TFeeUpd + TBasicUpd[update renameTo bbUpdate]
  + TPosUpd[bUpdate renameTo bbUpdate][update renameTo bUpdate]

```

Fig. 10. Class definitions for the different bank accounts.

```

class CFeeAccount implements IFeeAccount {
  int bal, nat fee;
  void deposit (nat x) {bal=bal+x;}
  void withdraw (nat x) {update(x);}
  void update (nat x) {bUpdate(x+fee);}
  void bUpdate (nat x) {bal=bal-x;}
}

```

Fig. 11. The flattened version of class CFeeAccount.

$\mathcal{C}_0$ and $\mathcal{T}_0$ where $\mathcal{C}_0$ is empty and where each basic trait is bound to an annotated basic trait expression in $\mathcal{T}_0$, corresponding to the definition found in Figure 8. The environment $\mathcal{T}_0$ is further detailed in Appendix C. Figure 12 shows the valid specifications for the methods in the classes CFeeAccount, CBasicAccount, CPosAccount, CPosFeeAccount, and CFeePosAccount. Furthermore, the trait in which the method is defined is indicated in the column *Trait*, and the column labelled *PO* indicates whether a new specification for the trait was needed to make the proof go through. Thus, if $\mathcal{C}$ and $\mathcal{T}$ are the class and trait environments resulting from the class analysis, a row $C \ m \ sp \ TAE$ in the table means that $sp \in S_{\mathcal{C}}(C, m)$ and $sp \in specs(\mathcal{T}(TAE)(m))$, and an $*$ indicates that a new specification for the method was provided during the analysis of C .

The specifications for each class in Figure 12 are consistent: Each internal call requirement follows from the guarantee of the called method, and each method contract in the interface follows from the corresponding guarantee in the class. If an interface has two contracts for *withdraw*, both are entailed by the verified specification. Note that the number of user-supplied specifications may be reduced by introducing uninterpreted boolean assertions in specifications (see [DDJ⁺12]). For instance, *withdraw* in TBasicAccount could then be specified by $\langle(p, q), update : (p, q)\rangle$ for some p and q , which means that it would suffice to provide only one proof outline for this method.

9. Related Work

This section relates the presented work on a proof system for a trait-based object-oriented language to related work on proof systems for object-oriented programming languages. We refer to [BDSS13] for an in-depth discussion of traits

Class	Method	Valid specifications	Trait	PO
CFeeAccount	withdraw	$\langle\langle bal == b_0, bal == b_0 - x - fee \rangle, \text{update} : (bal == b_0, bal == b_0 - x - fee) \rangle\rangle$	TBasicAccount	*
	update	$\langle\langle bal == b_0, bal == b_0 - x - fee \rangle, bUpdate : (bal == b_0, bal == b_0 - x) \rangle\rangle$	TFeeUpd	*
	bUpdate	$\langle\langle bal == b_0, bal == b_0 - x \rangle, \emptyset \rangle\rangle$	TBasicUpd[bU]	
CBasicAccount	withdraw	$\langle\langle bal == b_0, bal == b_0 - x \rangle, \text{update} : (bal == b_0, bal == b_0 - x) \rangle\rangle$	TBasicAccount	*
	update	$\langle\langle bal == b_0, bal == b_0 - x \rangle, \emptyset \rangle\rangle$	TBasicUpd	
CPosAccount	withdraw	$\langle\langle bal == b_0, (b_0 \geq x \Rightarrow bal == b_0 - x) \wedge (b_0 < x \Rightarrow bal == b_0) \rangle, \text{update} : (bal == b_0, (b_0 \geq x \Rightarrow bal == b_0 - x) \wedge (b_0 < x \Rightarrow bal == b_0)) \rangle\rangle$	TBasicAccount	*
	update	$\langle\langle bal == b_0, (b_0 \geq x \Rightarrow bal == b_0 - x) \wedge (b_0 < x \Rightarrow bal == b_0) \rangle, bUpdate : (bal == b_0, bal == b_0 - x) \rangle\rangle$	TPosAccount	*
	bUpdate	$\langle\langle bal == b_0, bal == b_0 - x \rangle, \emptyset \rangle\rangle$	TBasicUpd[bU]	
CPosFeeAccount	withdraw	$\langle\langle bal == b_0, (b_0 \geq x \Rightarrow bal == b_0 - x - fee) \wedge (b_0 < x \Rightarrow bal == b_0) \rangle, \text{update} : (bal == b_0, (b_0 \geq x \Rightarrow bal == b_0 - x - fee) \wedge (b_0 < x \Rightarrow bal == b_0)) \rangle\rangle$	TBasicAccount	*
	update	$\langle\langle bal == b_0, (b_0 \geq x \Rightarrow bal == b_0 - x - fee) \wedge (b_0 < x \Rightarrow bal == b_0) \rangle, bUpdate : (bal == b_0, bal == b_0 - x - fee) \rangle\rangle$	TPosAccount	*
	bUpdate	$\langle\langle bal == b_0, bal == b_0 - x - fee \rangle, bbUpdate : (bal == b_0, bal == b_0 - x) \rangle\rangle$	TFeeUpd[bbU]	
	bbUpdate	$\langle\langle bal == b_0, bal == b_0 - x \rangle, \emptyset \rangle\rangle$	TBasicUpd[bbU]	
CFeePosAccount	withdraw	$\langle\langle bal == b_0, (b_0 \geq (x + fee) \Rightarrow bal == b_0 - x - fee) \wedge (b_0 < (x + fee) \Rightarrow bal == b_0) \rangle, \text{update} : (bal == b_0, (b_0 \geq (x + fee) \Rightarrow bal == b_0 - x - fee) \wedge (b_0 < (x + fee) \Rightarrow bal == b_0)) \rangle\rangle$	TBasicAccount	*
	update	$\langle\langle bal == b_0, (b_0 \geq (x + fee) \Rightarrow bal == b_0 - x - fee) \wedge (b_0 < (x + fee) \Rightarrow bal == b_0) \rangle, bUpdate : (bal == b_0, (b_0 \geq x \Rightarrow bal == b_0 - x) \wedge (b_0 < x \Rightarrow bal == b_0)) \rangle\rangle$	TFeeAccount	*
	bUpdate	$\langle\langle bal == b_0, (b_0 \geq x \Rightarrow bal == b_0 - x) \wedge (b_0 < x \Rightarrow bal == b_0) \rangle, bbUpdate : (bal == b_0, bal == b_0 - x) \rangle\rangle$	TPosUpd[bbU]	
	bbUpdate	$\langle\langle bal == b_0, bal == b_0 - x \rangle, \emptyset \rangle\rangle$	TBasicUpd[bbU]	

Abbreviations for trait expressions:

$TBasicUpd[bU] = TBasicUpd[\text{update } \textbf{renameTo } bUpdate]$
 $TBasicUpd[bbU] = TBasicUpd[\text{update } \textbf{renameTo } bbUpdate]$
 $TFeeUpd[bbU] = TFeeUpd[bUpdate \textbf{ renameTo } bbUpdate] [\text{update } \textbf{renameTo } bUpdate]$
 $TPosUpd[bbU] = TPosUpd[bUpdate \textbf{ renameTo } bbUpdate] [\text{update } \textbf{renameTo } bUpdate]$

Fig. 12. Summary of the analysis of the `withdraw` method in each class.

and of how the trait language considered in this paper relates to other trait languages. We do not attempt to give a general discussion of object-oriented program verification here, but focus on how different approaches address code reuse mechanisms.

Multiple specifications of methods have been recognized as a convenient way of specifying behavior, for example in Fresco capsules [Wil91], with the also-constructs of JML [BCC⁺05], and in Parkinson and Bierman’s work [PB08]. In these approaches, two specifications are flattened (or “desugared”) into a single specification by a composition rule; e.g., in Fresco two specifications (p_1, q_1) and (p_2, q_2) of the same method have the interpretation $(p_1 \vee p_2, (p_1 \rightarrow q_1^o) \wedge (p_2 \rightarrow q_2^o))$ according to our notation. Note that this example demonstrates how it is never necessary to allow multiple specifications, but rather it is a convenience which is particularly clear with incremental approaches to deductive reasoning. In the context of the flexible code reuse offered by traits, there can be several specifications and we have therefore opted for an approach based on reasoning over sets of specifications instead of flattening these sets.

Single inheritance is the object-oriented code reuse mechanism which is by far the most studied and best supported in formal systems for program analysis. In the context of single inheritance, behavioral reasoning about extensible class hierarchies with late bound method calls is often performed in the context of *behavioral subtyping* (see, e.g., [LW94, Ame91, LN06, PHM99]). Behavioral subtyping is an incremental reasoning strategy in the sense that a subclass may be analyzed in the context of previously defined classes. In order to avoid reverification of superclasses, method overriding must preserve the specifications of overridden methods. This approach has also been

used for SCALA’s ‘trait’ construct, but “*significantly reduced the applicability and thereby benefits of traits*” [Sch10]. Combining separation logic with object-oriented structuring mechanisms, Parkinson and Bierman propose abstract predicate families which relate a set of implementations to the late binding mechanism [PB08]. The approach separates behavioral subtyping requirements on dynamic specifications from code reuse flexibility on static specifications and goes beyond behavioral subtyping by allowing new specifications of inherited code in a subclass, which need not respect the dynamic specifications of the superclass. *Lazy behavioral subtyping* [DJOS10] is an incremental reasoning strategy which supports more flexible code reuse than behavioral subtyping. With lazy behavioral subtyping, the *requirements* that a method guarantee imposes on late bound method calls are identified, and the main idea is that there is no need to preserve the full specifications of overridden methods. In order to avoid reverification of superclass methods, only the weaker *requirements* imposed on late bound method calls need to be preserved by method redefinitions in subclasses. Lazy behavioral subtyping is more flexible than the approach of [PB08] as it completely separates interface inheritance (for late bound external calls) and code reuse, and only requires compliance for requirements on internal calls instead of their specifications. Although traits do not support late bound internal calls, the flexible reuse of traits motivates us to follow the lazy behavioral subtyping approach in maintaining a separation of concerns between required and guaranteed assertions for method calls and definitions, respectively, and in allowing sets of specifications for each method definition.

Multiple inheritance is widely used in modeling notations such as UML [BRJ99], to capture that a concept naturally inherits from several other concepts. Versions of multiple inheritance are found in C++, CLOS, Eiffel, and OCaml. Creol [JOY06] has proposed a so-called healthy binding strategy which resolves horizontal name conflicts by avoiding accidental overriding. The proof systems presented in [NCMM09, DJOS11, LQ08, vSC10] are the only proof systems we know for multiple inheritance class hierarchies. The work in [NCMM09] presents a Hoare-style program logic for Eiffel that handles multiple inheritance based on an existing program logic for single inheritance by extending the method lookup definition. In [LQ08], method calls are assumed to be fully qualified in order to avoid ambiguities, and diamond inheritance is not considered. In [vSC10], ambiguities are assumed to be resolved by the programmer, a method can only be inherited if it has the same implementation in all parents. In contrast, [DJOS11] applies lazy behavioral subtyping to multiple inheritance and shows that healthy method binding is sufficient to allow incremental reasoning about multiple inheritance class hierarchies. The main challenges for reasoning about class hierarchies with multiple inheritance are related to late bound method calls, such as accidental or ambiguous overriding (sometimes called the “diamond problem”). In contrast, traits do not support late binding and require explicit disambiguation from the programmer. However, the flexible composition supported by traits necessitates a *delayed selection* of relevant method specifications, such that the requirements in our proof system are first checked when classes are assembled from traits. Technically, this makes lazy behavioral subtyping fairly different from the proposed proof system for traits. We are not aware of any previous proposal for a deductive proof system for a trait-based language.

Invasive code reuse mechanisms such as aspect-oriented programming (AOP) [KLM⁺97, FES10], feature-oriented programming (FOP) [BSR04, AKL13] and delta-oriented programming (DOP) [SBB⁺10] have been proposed to improve code reuse at different levels of the software design by allowing methods defined in one module to be intercepted or changed in another module. Similar to traits, this code reuse flexibility poses a challenge for program verification. In contrast to traits, AOP, FOP and DOP have a notion of original call, which allows the redefined method to be called from inside the redefinition. Most prominently, AOP aims to modularize code that crosscuts the basic program modules [KLM⁺97]; so-called advice is used to allow a definition in an aspect to affect methods defined elsewhere. The two levels of modularity supported by AOP often improve code reuse in a program, compared to a single level of modularity as in traditional object-oriented programming. However, the highly invasive nature of aspects makes it very challenging to reason formally about program behavior. In contrast to traits, aspects rarely modify the state of a program in practice [KF07]; i.e., the aspects are largely external to the functionality of the core programs. Neither Aldrich [Ald05] nor Krishnamurthi and Fisler [KF07] address aspects which modify the program state in their work. Extending AOP with open modules allow hiding implementation details from the advice mechanism, by using interfaces to export pointcuts. For a functional core calculus of open modules, Aldrich uses weak bisimulation techniques to reason about module equivalence with respect to the application of advice [Ald05]. Krishnamurthi and Fisler use CTL model checking to verify AOP programs expressed as state machines with pointcut interfaces [KF07], to identify when the application of an aspect may violate properties of the original program. Similar to our work their approach is incremental and uses a cache. Extending the idea of pointcut interfaces, translucent contracts have been proposed to behaviorally restrict the applicability of aspects [BRLM11], thereby making AOP amenable to modular deductive verification. These works on AOP are akin to ours in that they address reasoning about mechanisms which aim at better code reuse than class inheritance. However, the proposed use of interfaces between classes and aspects would quickly become unwieldy if adapted to the composition of trait expressions.

Whereas AOP addresses changes to method calls in terms of pointcuts, FOP and DOP address changes to a set

of classes to accommodate feature selection in software product lines, allowing in particular methods to be redefined. Verification of FOP programs based on a meta-representation of all program variants that can be generated from a set of feature modules is considered in [TSAH12]. In [TSAH11], proof scripts for single feature modules are extracted from proofs for complete program variants and later composed for other program variants using these feature modules. Hähnle and Schaefer propose an approach to deductive verification for DOP which relies on behavioral subtyping for DOP [HS12]; i.e., specifications of methods introduced by deltas must be more specific than previous versions of these methods. Thus, a delta can only modify state in a way which refines the program state and each delta can be verified by approximating called methods defined in other deltas by the specification of their first declaration. We have proposed a different approach [DDJ⁺12]; which considers deltas as transformations of a program, and delta specifications as higher-order specifications which are instantiated at the delta composition time to transform a specific specification. This transformational approach uses symbolic assumptions on called methods and thus separates the specifications of method implementations from the requirements to method calls in a way which is similar to lazy behavioral subtyping [DJOS10]. The transformational approach leads to a two-phase verification process: the verification of deltas and the verification of the actual products based on the specifications already established for the used deltas. We believe that a transformational approach similar to [DDJ⁺12] could be applied to traits and a comparison between the approach taken in this paper and a transformational proof system would be an interesting extension of the work presented in this paper.

10. Conclusion and Future Work

This paper describes an approach to behavioral reasoning about trait-based object-oriented programs. Traits have been proposed as a particularly flexible way to achieve a high degree of fine-grained code reuse. We develop a deductive proof system, PST(PL), which reflects this fine-grained reuse potential at the level of behavioral reasoning. The approach focusses on verifying interface contracts for classes assembled from traits, based on method specifications in the traits. For flexible reuse, methods in traits have associated sets of possible specifications. These sets can be extended with new specifications as needed for class verification in an incremental way. When verifying a class, the interface contracts of the class drive the selection of possible specifications from the specification sets in the traits, in such a way that the internal consistency among the selected specifications is guaranteed and such that the interface contracts are valid. We show the soundness of PST(PL).

Trait-based code reuse is an interesting challenge for deductive proof systems because a high degree of reuse in a code structuring mechanism typically limits the degree of reasoning reuse which can be supported by the proof system. In particular, a trait cannot be fully verified independently of its context of composition. The presented approach addresses this challenge by proposing an incremental and compositional proof system for traits, in which specifications for methods in trait alteration expressions and in composed trait expressions are compositionally derived from basic trait expressions, and in which the sets of possible specifications can be incremented at need during the verification process. The approach combines bottom-up compositional reasoning about trait expressions with top-down selection of consistent specifications for methods declared in different traits during the verification process for classes.

The complexity of verifying trait-based programs using our proof system can be compared to the naive approach of first flattening the programs and then verifying them. When flattening each program built from a set of traits before verifying it, we need to verify each method in isolation which requires as many proofs as there are methods. In this case, there is no specification reuse. In the completely modular case of our proof system where the specifications in the trait environment suffice to establish the external requirements, we do not need any new code analysis to verify the new program. In the fully incremental case of our proof system where each external requirement needs a new trait specification to be verified, we need one proof for each method which is the same as for the naive approach. In all cases where some external requirements can be shown modularly from existing specifications and some external specifications need incremental verification by adding new specifications, we have some specification reuse and thereby less code analysis than when verifying flattened programs. With our approach, the gain comes from analyzing many programs with the incrementally constructed trait environment, to achieve specification reuse similar to modular reasoning for common ways of using of traits.

In this paper, we have concentrated on verifying method contracts given as pre- and postconditions. As future work, it is interesting to investigate to what extent invariants could be useful for trait-based programs, both at the level of classes and traits. A trait invariant can, for instance, capture relations between the required fields of a trait; this will extend the range of properties that can be incrementally verified for trait-based programs. Further, we plan to extend the KeY system [BHS07] for deductive verification of JAVA programs to a trait-based language such as TRAITRECORDJ [BDSS13] and to implement the proof system proposed in this paper within KeY.

Acknowledgements. The authors are grateful to Wolfgang Ahrendt, Richard Bubel, Olaf Owe, and Volker Stolz for valuable discussions on the subject of this work, and to the FAOC editor and anonymous reviewers for many useful comments and suggestions for improving the presentation and for further pointers to related work.

References

- [AdBO09] Krzysztof R. Apt, Frank S. de Boer, and Ernst-Rüdiger Olderog. *Verification of Sequential and Concurrent Systems*. Texts and Monographs in Computer Science. Springer, 3rd edition, 2009.
- [AKL13] Sven Apel, Christian Kästner, and Christian Lengauer. Language-independent and automated software composition: The Feature-House experience. *IEEE Transactions on Software Engineering*, 39(1):63–79, 2013.
- [Ald05] Jonathan Aldrich. Open modules: Modular reasoning about advice. In Andrew P. Black, editor, *Proc. European Conference on Object-Oriented Programming (ECOOP)*, volume 3586 of *Lecture Notes in Computer Science*, pages 144–168. Springer, 2005.
- [ALZ03] Davide Ancona, Giovanni Lagorio, and Elena Zucca. Jam—designing a Java extension with mixins. *Transactions on Programming Languages and Systems*, 25(5):641–712, September 2003.
- [Ame91] Pierre America. Designing an object-oriented programming language with behavioural subtyping. In J. W. de Bakker, W.-P. de Roever, and G. Rozenberg, editors, *Foundations of Object-Oriented Languages*, volume 489 of *Lecture Notes in Computer Science*, pages 60–90. Springer, 1991.
- [Apt81] Krzysztof R. Apt. Ten years of Hoare’s logic: A survey — Part I. *Transactions on Programming Languages and Systems*, 3(4):431–483, October 1981.
- [BC90] Gilad Bracha and William Cook. Mixin-based inheritance. In Norman Meyrowitz, editor, *Proceedings of the Conference on Object-Oriented Programming: Systems, Languages, and Applications / European Conference on Object-Oriented Programming*, pages 303–311. ACM Press, 1990.
- [BCC⁺05] Lilian Burdy, Yoonsik Cheon, David R. Cok, Michael D. Ernst, Joseph R. Kiniry, Gary T. Leavens, K. Rustan M. Leino, and Erik Poll. An overview of JML tools and applications. *International Journal on Software Tools for Technology Transfer*, 7(3):212–232, 2005.
- [BCD12] Lorenzo Bettini, Sara Capecchi, and Ferruccio Damiani. On flexible dynamic trait replacement for Java-like languages. *Science of Computer Programming*, 2012. Available online doi:10.1016/j.scico.2012.11.003.
- [BDD⁺10] Lorenzo Bettini, Ferruccio Damiani, Marco De Luca, Kathrin Geilmann, and Jan Schäfer. A calculus for boxes and traits in a Java-like setting. In Dave Clarke and Gul Agha, editors, *Coordination Models and Languages*, volume 6116 of *Lecture Notes in Computer Science*, pages 46–60. Springer, 2010.
- [BDG07] Viviana Bono, Ferruccio Damiani, and Elena Giachino. Separating Type, Behavior, and State to Achieve Very Fine-grained Reuse. In *Proc. Formal Techniques for Java-like Programs (FTfJP)*, 2007.
- [BDG08] Viviana Bono, Ferruccio Damiani, and Elena Giachino. On Traits and Types in a Java-like setting. In *Fifth IFIP International Conference On Theoretical Computer Science (TCS’08)*, volume 273 of *International Federation for Information Processing*, pages 367–382. Springer, 2008.
- [BDGS13] Lorenzo Bettini, Ferruccio Damiani, Kathrin Geilmann, and Jan Schäfer. Combining traits with boxes and ownership types in a Java-like setting. *Science of Computer Programming*, 78(2):218–247, 2013.
- [BDN⁺09] Andrew Black, Stéphane Ducasse, Oscar Nierstrasz, Damien Pollet, Damien Cassou, and Marcus Denker. *Pharo by Example*. Square Bracket Associates, 2009.
- [BDNW08] Alexandre Bergel, Stéphane Ducasse, Oscar Nierstrasz, and Roel Wuyts. Stateful traits and their formalization. *Computer Languages, Systems & Structures*, 34(2-3):83–108, 2008.
- [BDS09] Lorenzo Bettini, Ferruccio Damiani, and Ina Schaefer. Implementing SPL using Traits. Technical report, Dipartimento di Informatica, Università di Torino, 2009. Available at www.di.unito.it/~damiani/papers/isplurat.pdf.
- [BDS10] Lorenzo Bettini, Ferruccio Damiani, and Ina Schaefer. Implementing software product lines using traits. In *Proceedings of the 2010 ACM Symposium on Applied Computing (SAC’10)*, pages 2096–2102. ACM Press, 2010.
- [BDSS10] Lorenzo Bettini, Ferruccio Damiani, Ina Schaefer, and Fabio Strocchio. A prototypical Java-like language with records and traits. In *Proceedings of the 8th International Conference on the Principles and Practice of Programming in Java (PPPJ’10)*, pages 129–138. ACM Press, 2010.
- [BDSS13] Lorenzo Bettini, Ferruccio Damiani, Ina Schaefer, and Fabio Strocchio. TRAITRECORDJ: A programming language with traits and records. *Science of Computer Programming*, 2013. In press, available online doi:10.1016/j.scico.2011.06.007.
- [BHS07] Bernhard Beckert, Reiner Hähnle, and Peter H. Schmitt, editors. *Verification of Object-Oriented Software: The KeY Approach*, volume 4334 of *Lecture Notes in Computer Science*. Springer, 2007.
- [BRJ99] Grady Booch, James E. Rumbaugh, and Ivar Jacobson. *The Unified Modeling Language User Guide*. Addison-Wesley, 1999.
- [BRLM11] Mehdi Bagherzadeh, Hridesh Rajan, Gary T. Leavens, and Sean L. Mooney. Translucid contracts: expressive specification and modular verification for aspect-oriented interfaces. In Paulo Borba and Shigeru Chiba, editors, *Proceedings of the 10th International Conference on Aspect-Oriented Software Development (AOSD 2011)*, pages 141–152. ACM Press, 2011.
- [BSR04] Don Batory, Jacob Neal Sarvela, and Axel Rauschmayer. Scaling step-wise refinement. *IEEE Transactions on Software Engineering*, 30:355–371, 2004.
- [BvW98] Ralph-Johan Back and Joakim von Wright. *Refinement Calculus: A Systematic Introduction*. Graduate Texts in Computer Science. Springer, 1998.
- [dB99] Frank S. de Boer. A WP-calculus for OO. In Wolfgang Thomas, editor, *Proceedings of Foundations of Software Science and Computation Structure (FOSSACS’99)*, volume 1578 of *Lecture Notes in Computer Science*, pages 135–149. Springer, 1999.
- [DDJ⁺12] Ferruccio Damiani, Johan Dovland, Einar Broch Johnsen, Olaf Owe, Ina Schaefer, and Ingrid Chieh Yu. A transformational proof system for delta-oriented programming. In Eduardo Santana de Almeida, Christa Schwanninger, and David Benavides, editors, *Proc. 16th International Software Product Line Conference (SPLC’12)*, Volume 2, pages 53–60. ACM Press, 2012.

- [DDJS11] Ferruccio Damiani, Johan Dovland, Einar Broch Johnsen, and Ina Schaefer. Verifying traits: A proof system for fine-grained reuse. In *Proc. Formal Techniques for Java-like Programs (FTJP)*, pages 8:1–8:6. ACM Press, 2011.
- [DJOS10] Johan Dovland, Einar Broch Johnsen, Olaf Owe, and Martin Steffen. Lazy behavioral subtyping. *Journal of Logic and Algebraic Programming*, 79(7):578–607, 2010.
- [DJOS11] Johan Dovland, Einar Broch Johnsen, Olaf Owe, and Martin Steffen. Incremental reasoning with lazy behavioral subtyping for multiple inheritance. *Science of Computer Programming*, 76(10):915–941, 2011.
- [DNS⁺06] Stéphane Ducasse, Oscar Nierstrasz, Nathanael Schärli, Roel Wuyts, and Andrew P. Black. Traits: A mechanism for fine-grained reuse. *Transactions on Programming Languages and Systems*, 28(2):331–388, 2006.
- [FES10] Bruno De Fraine, Erik Ernst, and Mario Südholt. Essential AOP: The A calculus. In Theo D’Hondt, editor, *Proc. European Conference on Object-Oriented Programming (ECOOP)*, volume 6183 of *Lecture Notes in Computer Science*, pages 101–125. Springer, 2010.
- [FKF98] Matthew Flatt, Shriram Krishnamurthi, and Matthias Felleisen. Classes and mixins. In *Proc. Principles of Programming Languages (POPL)*, pages 171–183. ACM Press, 1998.
- [HLL⁺12] John Hatcliff, Gary T. Leavens, K. Rustan M. Leino, Peter Müller, and Matthew J. Parkinson. Behavioral interface specification languages. *ACM Computing Surveys*, 44(3):16, 2012.
- [Hoa69] Charles Antony Richard Hoare. An Axiomatic Basis of Computer Programming. *Communications of the ACM*, 12:576–580, 1969.
- [Hoa71] Charles Antony Richard Hoare. Procedures and parameters: An axiomatic approach. In E. Engeler, editor, *Symposium On Semantics of Algorithmic Languages*, volume 188 of *Lecture Notes in Mathematics*, pages 102–116. Springer, 1971.
- [HS12] Reiner Hähnle and Ina Schaefer. A Liskov principle for delta-oriented programming. In Tiziana Margaria and Bernhard Steffen, editors, *Proc. 5th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA)*, volume 7609 of *Lecture Notes in Computer Science*, pages 32–46. Springer, 2012.
- [IPW01] Atsushi Igarashi, Benjamin C. Pierce, and Philip Wadler. Featherweight Java: A minimal core calculus for Java and GJ. *Transactions on Programming Languages and Systems*, 23(3):396–450, 2001.
- [JOY06] Einar Broch Johnsen, Olaf Owe, and Ingrid Chieh Yu. Creol: A type-safe object-oriented model for distributed concurrent systems. *Theoretical Computer Science*, 365(1–2):23–66, November 2006.
- [KF07] Shriram Krishnamurthi and Kathi Fisler. Foundations of incremental aspect model-checking. *Transactions on Software Engineering and Methodology*, 16(2), 2007.
- [KLM⁺97] Gregor Kiczales, John Lamping, Anurag Mendhekar, Chris Maeda, Cristina Lopes, Jean-Marc Loingtier, and John Irwin. Aspect-oriented programming. In Mehmet Akşit and Satoshi Matsuoka, editors, *Proc. European Conference on Object-Oriented Programming (ECOOP)*, volume 1241 of *Lecture Notes in Computer Science*, pages 220–242. Springer, 1997.
- [LM96] Marc Van Limberghen and Tom Mens. Encapsulation and composition as orthogonal operators on mixins: A solution to multiple inheritance problems. *Object Oriented Systems*, 3(1):1–30, 1996.
- [LN06] Gary T. Leavens and David A. Naumann. Behavioral subtyping, specification inheritance, and modular reasoning. Technical Report 06-20a, Department of Computer Science, Iowa State University, Ames, Iowa, 2006.
- [LQ08] Chenguang Luo and Shengchao Qin. Separation logic for multiple inheritance. *Electronic Notes in Theoretical Computer Science*, 212:27–40, April 2008. Proc. First Intl. Conf. on Foundations of Informatics, Computing and Software (FICS).
- [LS08a] Luigi Liquori and Arnaud Spiwack. Extending FeatherTrait Java with interfaces. *Theoretical Computer Science*, 398(1-3):243–260, 2008.
- [LS08b] Luigi Liquori and Arnaud Spiwack. FeatherTrait: A modest extension of Featherweight Java. *Transactions on Programming Languages and Systems*, 30(2):1–32, 2008.
- [LSZ09] Giovanni Lagorio, Marco Servetto, and Elena Zucca. Flattening versus direct semantics for Featherweight Jigsaw. In *Proc. Workshop on Foundations of Object-Oriented Languages (FOOL)*, 2009.
- [LW94] Barbara Liskov and Jeannette M. Wing. A behavioral notion of subtyping. *Transactions on Programming Languages and Systems*, 16(6):1811–1841, November 1994.
- [NCMM09] Martin Nordio, Cristiano Calcagno, Peter Müller, and Bertrand Meyer. A Sound and Complete Program Logic for Eiffel. In M. Oriol, editor, *TOOLS-EUROPE 2009*, volume 33 of *Lecture Notes in Business and Information Processing*, pages 195–214, 2009.
- [NDS06] Oscar Nierstrasz, Stéphane Ducasse, and Nathanael Schärli. Flattening traits. *Journal of Object Technology*, 5(4):129–148, 2006.
- [OG76] Susan Owicki and David Gries. An axiomatic proof technique for parallel programs I. *Acta Informatica*, 6(4):319–340, 1976.
- [OSV10] Martin Odersky, Lex Spoon, and Bill Venners. *Programming in Scala*. Artima Press, 2 edition, 2010.
- [PB08] Matthew J. Parkinson and Gavin M. Bierman. Separation logic, abstraction and inheritance. In *Proc. Principles of Programming Languages (POPL)*, pages 75–86. ACM Press, 2008.
- [PBC06] Matthew J. Parkinson, Richard Bornat, and Cristiano Calcagno. Variables as resource in Hoare logics. In *Proc. Symposium on Logic in Computer Science (LICS’06)*, pages 137–146. IEEE Computer Society Press, 2006.
- [PHM99] Arnd Poetzsch-Heffter and Peter Müller. A programming logic for sequential Java. In S. Doaitse Swierstra, editor, *8th European Symposium on Programming Languages and Systems (ESOP’99)*, volume 1576 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1999.
- [RT06] John H. Reppy and Aaron Turon. A foundation for trait-based metaprogramming. In *Proceedings of FOOL/WOOD*, 2006.
- [RT07] John H. Reppy and Aaron Turon. Metaprogramming with traits. In *Proc. European Conference on Object-Oriented Programming (ECOOP)*, volume 4609 of *Lecture Notes in Computer Science*, pages 373–398. Springer, 2007.
- [SAC⁺11] Guy L. Steele Jr., Eric E. Allen, David Chase, Christine H. Flood, Victor Luchangco, Jan-Willem Maessen, and Sukyoung Ryu. Fortress (Sun HPCS language). In David A. Padua, editor, *Encyclopedia of Parallel Computing*, pages 718–735. Springer, 2011.
- [SBB⁺10] Ina Schaefer, Lorenzo Bettini, Viviana Bono, Ferruccio Damiani, and Nico Tanzarella. Delta-Oriented Programming of Software Product Lines. In Jan Bosch and Jaejoon Lee, editors, *Software Product Lines: Going Beyond (SPLC’10)*, volume 6287 of *Lecture Notes in Computer Science*, pages 77–91. Springer, 2010.
- [Sch10] Malte Schwerhoff. Verifying Scala traits. Semester Report, Swiss Federal Institute of Technology Zurich (ETH), October 2010.
- [SD05] Charles Smith and Sophia Drossopoulou. Chai: Traits for Java-like languages. In *Proc. European Conference on Object-Oriented Programming (ECOOP)*, volume 3586 of *Lecture Notes in Computer Science*, pages 453–478. Springer, 2005.

$$\begin{aligned}
\mathcal{T} \oplus (\text{TAE}, \text{ao}) &\triangleq \mathcal{T}[\text{TAE} \text{ ao} \mapsto \text{mod}(\mathcal{T}(\text{TAE}), \text{ao})] \\
\text{mod}(\{\bar{F}; \bar{H}; \bar{AM}\}, [\text{exclude } m]) &\triangleq \begin{cases} \{\bar{F}; \bar{H}; \bar{AM}\} \\ \{\text{fields}(\text{remMtd}(\bar{AM}, m)); \text{calls}(\text{remMtd}(\bar{AM}, m)); \text{remMtd}(\bar{AM}, m)\} \end{cases} \quad \begin{array}{l} \text{if } m \notin \bar{AM} \\ \text{otherwise} \end{array} \\
\text{mod}(\{\bar{F}; \bar{H}; \bar{AM}\}, [m \text{ aliasAs } m']) &\triangleq \begin{cases} \{\bar{F}; \bar{H}; \bar{AM}\} \\ \{\bar{F}; \bar{H} \setminus m'; \bar{AM} \text{ renId}(\bar{AM}(m), m, m')\} \end{cases} \quad \begin{array}{l} \text{if } m \notin \bar{AM} \\ \text{otherwise} \end{array} \\
\text{mod}(\{\bar{F}; \bar{H}; \bar{AM}\}, [m \text{ renameTo } m']) &\triangleq \{\bar{F}; \text{calls}(\text{renMtd}(\bar{AM}, m, m')); \text{renMtd}(\bar{AM}, m, m')\} \\
\text{mod}(\{\bar{F}; \bar{H}; \bar{AM}\}, [f \text{ renameTo } f']) &\triangleq \{\text{fields}(\text{renFld}(\bar{AM}, f, f')); \bar{H}; \text{renFld}(\bar{AM}, f, f')\} \\
\text{remMtd}(\emptyset, m) &\triangleq \emptyset \\
\text{remMtd}(\text{I } n(\bar{I} \bar{x})\{t\} \bar{sp} \bar{AM}, m) &\triangleq \begin{cases} \bar{AM} \\ \text{I } n(\bar{I} \bar{x})\{t\} \bar{sp} \text{remMtd}(\bar{AM}, m) \end{cases} \quad \begin{array}{l} \text{if } m = n \\ \text{otherwise} \end{array} \\
\text{renId}(\text{I } n(\bar{I} \bar{x})\{t\} \bar{sp}, m, m') &\triangleq \text{I } n[m'/m](\bar{I} \bar{x})\{t\} \bar{sp} \\
\text{renMtd}(\emptyset, m, m') &\triangleq \emptyset \\
\text{renMtd}(\bar{AM} \bar{AM}, m, m') &\triangleq \text{renMtd}(\bar{AM}, m, m') \text{ renMtd}(\bar{AM}, m, m') \\
\text{renMtd}(\text{I } n(\bar{I} \bar{x})\{t\} \bar{sp}, m, m') &\triangleq \text{I } n[m'/m](\bar{I} \bar{x})\{t[m'/m]\} \text{renSM}(\bar{sp}, m, m') \\
\text{renFld}(\emptyset, f, f') &\triangleq \emptyset \\
\text{renFld}(\bar{AM} \bar{AM}, f, f') &\triangleq \text{renFld}(\bar{AM}, f, f') \text{ renFld}(\bar{AM}, f, f') \\
\text{renFld}(M \bar{sp}, f, f') &\triangleq \begin{cases} \text{renMF}(M, f, f') \text{ discard}(\bar{sp}, \{f, f', \text{assign}(M)\}) \\ \text{renMF}(M, f, f') \text{ renSF}(\text{discard}(\bar{sp}, \{f'\}), f, f') \\ M \text{ discard}(\bar{sp}, \{f\}) \end{cases} \quad \begin{array}{l} \text{if } f \in \text{fields}(M) \wedge f' \in \text{fields}(M) \\ \text{if } f' \notin \text{fields}(M) \\ \text{otherwise} \end{array} \\
\text{renSM}(\varepsilon, m, m') &\triangleq \varepsilon \\
\text{renSM}(\langle (p, q), \bar{R} \rangle \bar{sp}, m, m') &\triangleq \langle (p, q), \text{renRM}(\bar{R}, m, m') \rangle \text{renSM}(\bar{sp}, m, m') \\
\text{renRM}(\emptyset, m, m') &\triangleq \emptyset \\
\text{renRM}(n : (r, s) \bar{R}, m, m') &\triangleq n[m'/m] : (r, s) \text{renRM}(\bar{R}, m, m') \\
\text{renSF}(\varepsilon, f, f') &\triangleq \varepsilon \\
\text{renSF}(\langle (p, q), \bar{R} \rangle \bar{sp}, f, f') &\triangleq \langle (p[f'/f], q[f'/f]), \text{renRF}(\bar{R}, f, f') \rangle \text{renSF}(\bar{sp}, f, f') \\
\text{renRF}(\emptyset, f, f') &\triangleq \emptyset \\
\text{renRF}(n : (r, s) \bar{R}, f, f') &\triangleq n : (r[f'/f], s[f'/f]) \text{renRF}(\bar{R}, f, f') \\
\text{renMF}(\text{I } n(\bar{I} \bar{x})\{t\}, f, f') &\triangleq \text{I } n(\bar{I} \bar{x})\{t[f'/f]\} \\
\text{discard}(\varepsilon, \bar{f}) &\triangleq \varepsilon \\
\text{discard}(\text{sp } \bar{sp}, \bar{f}) &\triangleq \begin{cases} \text{discard}(\bar{sp}, \bar{f}) & \text{if } \bar{f} \cap \text{vars}(\text{sp}) \neq \emptyset \\ \text{sp } \text{discard}(\bar{sp}, \bar{f}) & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 13. The definition of $\mathcal{T} \oplus (\text{TAE}, \text{ao})$ with associated auxiliary functions.

- [SDNB03] Nathanael Schärli, Stéphane Ducasse, Oscar Nierstrasz, and Andrew Black. Traits: Composable units of behavior. In *Proc. European Conference on Object-Oriented Programming (ECOOP)*, volume 2743 of *Lecture Notes in Computer Science*, pages 248–274. Springer, 2003.
- [Tra11] TraitRecordJ website, May 2011. <http://traitrecordj.sourceforge.net/>.
- [TSAH12] Thomas Thüm, Ina Schaefer, Sven Apel, and Martin Hentschel. Family-based deductive verification of software product lines. In *International Conference on Generative Programming and Component Engineering (GPCE)*, pages 11–20. ACM Press, 2012.
- [TSKA11] Thomas Thüm, Ina Schaefer, Martin Kuhlemann, and Sven Apel. Proof composition for deductive verification of software product lines. In *Proc. Intl. Workshop on Variability-intensive Systems Testing, Validation and Verification*, pages 270–277. IEEE Computer Society Press, 2011.
- [vSC10] Stephan van Staden and Cristiano Calcagno. Reasoning about multiple related abstractions with MultiStar. In *Proc. Object-Oriented Programming Systems, Languages and Applications (OOPSLA)*, pages 504–519. ACM Press, 2010.
- [Wil91] Alan Wills. Capsules and types in Fresco: Program verification in Smalltalk. In Pierre America, editor, *Proc. European Conference on Object-Oriented Programming (ECOOP)*, volume 512 of *Lecture Notes in Computer Science*, pages 59–76. Springer, 1991.

A. Auxiliary Functions for Extending the Trait Environment

This section details the auxiliary functions used to update the trait environment for trait alteration expressions in the rule **EXTEND** of the proof system for trait analysis given in Section 7.2. The auxiliary functions are defined in Figure 13. The update function $\mathcal{T} \oplus (\text{TAE}, \text{ao})$ creates an entry for TAE ao in $\mathcal{T}$ by *modifying* the existing entry for TAE. For a set $\bar{M}$ of method definitions, let the functions $\text{fields}(\bar{M})$ and $\text{calls}(\bar{M})$ return the required fields and method signatures of $\bar{M}$, respectively. These functions are straightforwardly lifted to annotated methods by ignoring specifications (e.g., $\text{fields}(\bar{M} \ \overline{\text{sp}}) \triangleq \text{fields}(\bar{M})$) and to sets of annotated methods. For a set $\bar{H}$ of method signatures, let $\bar{H} \setminus m$ return the set $\bar{H}$ except the signature for m . Without loss of generality, we may assume that there are no name clashes between fields and locally defined variables in methods. Let $\text{assign}(\bar{M})$ denote the subset of $\text{fields}(\bar{M})$ to which a method M assigns values (the write-set of $\bar{M}$). Given a method specification sp , let $\text{vars}(\text{sp})$ denote the set of free variables in sp .

Trait alteration operations construct a new trait expression from an old trait expression. The modification function mod is defined by cases for the different trait alteration operations. For the operation $[\text{exclude } m]$, the new trait expression is obtained by removing the definition and specifications of m , if present, from the old trait expression. For the operation $[m \ \text{aliasAs } m']$, the new trait expression is obtained by duplicating the definition of m , if defined, and giving the copy the new name m' . The new method m' has the same specifications as the old method. For the $[m \ \text{renameTo } m']$ operation, the new trait expression is obtained by a consistent renaming in all method definitions and requirements.

As discussed in Section 5, field renaming requires special care since a direct substitution may invalidate specifications. For the operation $[\bar{f} \ \text{renameTo } \bar{f}']$, this is reflected in the definition of $\text{renFld}(\bar{M} \ \overline{\text{sp}}, \bar{f}, \bar{f}')$ which applies to each method in the old trait expression. If both $\bar{f}$ and $\bar{f}'$ occur in the method body, the new method is obtained by *discarding* all previous specifications, except those that do not contain variables occurring in $\bar{f}$, $\bar{f}'$, and the write-set of $\bar{M}$. Thus, the remaining specifications are neither affected by the substitution nor by the execution of $\bar{M}$. This is a strong restriction on the specifications, capturing that the behavior of the method has changed. Note that it is not a sufficient requirement that the specification itself does not contain $\bar{f}$ or $\bar{f}'$ in order to be preserved; for example $\{\text{true}\} \ \bar{f} := 5; \bar{f}' := 6; u := \bar{f} \ \{u == 5\}$ becomes invalid if $\bar{f}$ is replaced by $\bar{f}'$. The second line in the definition of $\text{renFld}(\bar{M} \ \overline{\text{sp}}, \bar{f}, \bar{f}')$ covers the case where the new field $\bar{f}'$ does not occur in the method body. In this case, specifications which contain the new field $\bar{f}'$ are discarded to avoid name clashes. In the remaining specifications, all occurrences of $\bar{f}$ are renamed. The **otherwise** case covers the situation $\bar{f}' \in \text{fields}(\bar{M}) \wedge \bar{f} \notin \text{fields}(\bar{M})$. In this case, the old field $\bar{f}$ does not occur in the method, which means that the method definition itself is not changed. To avoid name clashes, we then discard the specifications which contain $\bar{f}$. The soundness of $\mathcal{T} \oplus (\text{TAE}, \text{ao})$ follows by Lemma B.2. (A more fine-grained analysis of field renaming may be possible in program logics PL which support variables-as-resources [PBC06].)

The auxiliary function $\text{renSM}(\overline{\text{sp}}, m, m')$ renames the method m in $\overline{\text{sp}}$. Since method names do not occur in assertions, this function is defined in terms of renRM which replaces each requirement $n : (r, s)$ by $n[m'/m] : (r, s)$. Similarly, the function $\text{renSF}(\overline{\text{sp}}, \bar{f}, \bar{f}')$ renames the field $\bar{f}$ in $\overline{\text{sp}}$, and is defined in terms of renRF which renames fields in requirements. The function renMF renames fields in method definitions, and the function $\text{discard}(\overline{\text{sp}}, \bar{f})$ returns $\overline{\text{sp}}$ except specifications that contain variables in $\bar{f}$.

B. Soundness of PST(PL)

Appendix B.1 presents the proof for the soundness of PST(PL). The main soundness result is formalized by Theorem 7.4, which expresses that the verification of a class leads to a consistent set of specifications for that class. Appendix B.2 contains auxiliary lemmas used for the soundness proof. Especially, the proof of Lemma B.2 ensures that the trait environment $\mathcal{T}$ is sound after the successful analysis of the basic trait declarations; i.e., there is a valid proof outline for all specifications recorded in $\mathcal{T}$ during the analysis of traits and classes. We here assume that trait and class environments have been constructed from the empty trait and class environments by application of the rules of PST(PL).

B.1. Main Theorem

Theorem 7.4. *For a given class `class C implements I by {  $\overline{r\overline{p}_I, \overline{F}};$  } and CTE, if the successful analysis of C in PST(PL) leads to a class environment  $\mathcal{C}$ , then the set of method specifications for C in  $\mathcal{C}$  are consistent for the flattened version of C.`*

Proof. We first check Condition 2 in Definition 7.3. For each $m \in mtds(CTE)$, we consider the set $S_{\mathcal{C}}(C, m)$. Let $CTE \triangleq TAE_1 + \dots + TAE_n$. By type safety, there must be exactly one i (for $1 \leq i \leq n$) such that $m \in mtds(\mathcal{T}(TAE_i))$. By Lemma B.6 we know that $S_{\mathcal{C}}(C, m) \subseteq specs(\mathcal{T}(TAE_i)(m))$. Let $\langle (p, q), \overline{R} \rangle \in S_{\mathcal{C}}(C, m)$. By Lemma B.2, there exists a proof outline O for the method body of m with guarantee (p, q) such that $\overline{R} \rightarrow reqs(O)$. For each $n : (r, s) \in \overline{R}$, we prove $guar(S_{\mathcal{C}}(C, n)) \rightarrow (r, s)$ by induction over the inference rules of PST(PL).

To include $\langle (p, q), \overline{R} \rangle$ in $S_{\mathcal{C}}(C, m)$, rule ANALYZE must be applied, which again generates a *discharge* operation for a set of requirements which includes $\overline{R}$. Repeatedly applying rule DECOMP4 results in a *discharge*($n : (r, s)$) operation, which is analyzed by either OPENANALYSIS or CLOSEANALYSIS. If OPENANALYSIS is applied, we get an operation *analyze*($TAE_j, n : (r, s)$) where j is such that $n \in mtds(\mathcal{T}(TAE_j))$. This operation must again be handled by the ANALYZE rule, which ensures the entailment $guar(S_{\mathcal{C}}(C, n)) \rightarrow (r, s)$. Before ANALYZE can be applied in the proof, INCREMENT may be needed to introduce the necessary specifications $\overline{sp}$ into the trait environment to ensure that $guar(\overline{sp}) \rightarrow (r, s)$. The application of ANALYZE again generates a *discharge* operation for a set of requirements which includes $req(\overline{sp})$. Since the analysis of C succeeds, the repeated application of ANALYZE must eventually terminate. If CLOSEANALYSIS is applied, the conclusion $guar(S_{\mathcal{C}}(C, n)) \rightarrow (r, s)$ follows directly.

For Condition 1 in Definition 7.3, we have by the rules CLASS and DECOMP4 that a *discharge*($m : (r, s)$) operation is analyzed for each $m : (r, s) \in contracts(\overline{r\overline{p}_I}, \overline{I})$. By an argument corresponding to the one for *discharge* above, we get $S_{\mathcal{C}}(C, m) \rightarrow (r, s)$.

It follows from Lemma B.5 that the sets of method specifications $S_{\mathcal{C}}(C, m)$ for all $m \in mtds(CTE)$ constitute a consistent set of method specifications for the flattened version of C. $\square$

B.2. Auxiliary Lemmas

Lemma B.1. Let the trait alteration expression TAE be defined by $Tb \overline{a\overline{o}}$. If $TAE \in \mathcal{T}$, then for any (possibly empty) prefix $\overline{a\overline{o}'}$ of $\overline{a\overline{o}}$, we also have $Tb \overline{a\overline{o}'} \in \mathcal{T}$.

Proof. The proof is by induction over the length of the list $\overline{a\overline{o}}$ of alteration operations. If $\overline{a\overline{o}}$ is empty, then TAE is a basic trait Tb which is introduced in $\mathcal{T}$ by rule BASICTRAIT. For the induction step, consider $TAE \overline{a\overline{o}}$, where $TAE = Tb \overline{a\overline{o}}$. Since $TAE \overline{a\overline{o}}$ must have been introduced by rule EXTEND, we have that $TAE \in \mathcal{T}$, i.e., we have $Tb \overline{a\overline{o}} \in \mathcal{T}$. By the induction hypothesis, $Tb \overline{a\overline{o}'} \in \mathcal{T}$ for all prefixes $\overline{a\overline{o}'}$ of $\overline{a\overline{o}}$. $\square$

Lemma B.2 (Soundness of $\mathcal{T}$). Let $\mathcal{T}$ be a trait environment such that all basic traits in $\mathcal{T}$ are successfully analyzed. Then the following holds for all $TAE \in \mathcal{T}$:

$$\forall (M \overline{sp}) \in \mathcal{T}(TAE) . \forall \langle (p, q), \overline{R} \rangle \in \overline{sp} . \\ \exists O . O \vdash_{PL} body(M) : (p, q) \wedge \overline{R} \rightarrow reqs(O)$$

Proof. The proof is by induction over the structure of TAE. By Lemma B.1 each prefix of TAE is also in $\mathcal{T}$.

Base case: $TAE = Tb$. Then $\langle (p, q), \overline{R} \rangle$ is included in $\overline{sp}$ by either VERIFY or BASICTRAIT. If rule VERIFY is applied, the lemma follows immediately from the premise of the rule. The application of rule BASICTRAIT, possibly followed by a number of DECOMP1 and DECOMP2 applications, results in an operation *verify*($Tb, M \langle (p, q), \overline{R} \rangle$) for every specification $\langle (p, q), \overline{R} \rangle$ of some M in Tb . This operation is analyzed by VERIFY after zero or more applications of ADAPTATION. We close this case of the proof by induction over the number of applications of ADAPTATION. For zero applications, the lemma follows directly from the application of VERIFY. For the induction step, we assume a specification $\langle (p', q'), \overline{R}' \rangle$ such that $\langle (p', q'), \overline{R}' \rangle \rightarrow \langle (p, q), \overline{R} \rangle$ and $O' \vdash_{PL} body(M) : (p', q')$ and $\overline{R}' \rightarrow reqs(O')$ for some proof outline O' . By Lemma 4.7, there is a proof outline O such that $O \vdash_{PL} body(M) : (p, q)$ and $\overline{R} \rightarrow reqs(O)$.

Induction step: $TAE = TAE' \overline{a\overline{o}}$. By Lemma B.1, $TAE' \in \mathcal{T}$. Let $\mathcal{T}(TAE') = \{\overline{F}; \overline{H}; \overline{AM}\}$. The induction hypothesis IH assumes proof outlines for these methods; i.e., $\forall (M \overline{sp}) \in \overline{AM} . \forall \langle (p, q), \overline{R} \rangle \in \overline{sp} . \exists O . O \vdash_{PL} body(M) : (p, q) \wedge \overline{R} \rightarrow reqs(O)$.

The only rules that extend $\mathcal{T}$ with a new specification for non-basic traits are VERIFY and EXTEND. If $\langle (p, q), \overline{R} \rangle$ is included in the trait environment by VERIFY, the conclusion follows immediately. If EXTEND is applied, the trait

environment is extended to $\mathcal{T}[\text{TAE} \mapsto \text{mod}(\mathcal{T}(\text{TAE}'), \text{ao})]$, where the definition of mod can be found in Figure 13. The different cases for ao ensure the conditions of the lemma as follows:

- $\text{ao} = [\text{exclude } m]$. Here $\mathcal{T}(\text{TAE})$ contains the set $\overline{\text{AM}}' = \overline{\text{AM}} \setminus \overline{\text{AM}}(m)$. Proof outlines for $\overline{\text{AM}}'$ follow from *IH*.
- $\text{ao} = [m \text{ aliasAs } m']$. Then $\mathcal{T}(\text{TAE})$ contains the annotated methods $\overline{\text{AM}}$, and proof outlines for these specifications follow from *IH*. The trait $\mathcal{T}(\text{TAE})$ contains additional specifications if $m \in \overline{\text{AM}}$. Let $\overline{\text{AM}}(m) = \text{I } m(\overline{\text{I } \overline{x}})\{t\} \overline{\text{sp}}$. Then $\text{I } m'(\overline{\text{I } \overline{x}})\{t\} \overline{\text{sp}}$ is included in $\mathcal{T}(\text{TAE})$. Proof outlines for $\overline{\text{sp}}$ follow directly from *IH*.
- $\text{ao} = [m \text{ renameTo } m']$. For each $\text{I } n(\overline{\text{I } \overline{x}})\{t\} \overline{\text{sp}}$ in $\overline{\text{AM}}$, we have by Figure 13 that the following method is included in $\mathcal{T}(\text{TAE})$: $\text{I } n[m'/m](\overline{\text{I } \overline{x}})\{t[m'/m]\} \text{renSM}(\overline{\text{sp}}, m, m')$. For each $\langle (p, q), \overline{R} \rangle \in \overline{\text{sp}}$, we have by *IH* that there is a proof outline O for t such that $O \vdash_{PL} t : (p, q)$ where $\overline{R} \rightarrow \text{reqs}(O)$. From the definition of renSM in Figure 13, the specification $\langle (p, q), \text{renRM}(\overline{R}, m, m') \rangle$ is included in $\mathcal{T}(\text{TAE})(n[m'/m])$. From the definition of renRM we observe that if $\overline{R} \rightarrow \text{reqs}(O)$, then $\text{renRM}(\overline{R}, m, m') \rightarrow \text{renRM}(\text{reqs}(O), m, m')$. We construct a proof outline O' for $t[m'/m]$ by replacing each call $\{r\} m \{s\}$ in O by $\{r\} m' \{s\}$, i.e., $\text{reqs}(O') = \text{renRM}(\text{reqs}(O), m, m')$. The desired conclusion $\text{renRM}(\overline{R}, m, m') \rightarrow \text{reqs}(O')$ then follows since $\overline{R} \rightarrow \text{reqs}(O)$.
- $\text{ao} = [f \text{ renameTo } f']$. For each method $\text{I } m(\overline{\text{I } \overline{x}})\{t\} \overline{\text{sp}}$ in $\overline{\text{AM}}$, the modified method $\text{renFld}(\text{I } m(\overline{\text{I } \overline{x}})\{t\} \overline{\text{sp}}, f, f')$ is included in $\mathcal{T}(\text{TAE})$. The different cases in the definition of this function are considered separately.
 - Case $f \in \text{fields}(t) \wedge f' \in \text{fields}(t)$. In this case, f is replaced by f' in the method body t and all specifications which contain f , f' , or variables in $\text{assign}(t)$ are discarded. Thus, if a specification $\langle (p, q), \overline{R} \rangle$ is included, we have $\text{vars}(\langle (p, q), \overline{R} \rangle) \cap \{f, f', \text{assign}(t)\} = \emptyset$. By *IH*, there is a proof outline O such that $O \vdash_{PL} t : (p, q)$ where $\overline{R} \rightarrow \text{reqs}(O)$. It then follows that O is a proof outline for $t[f'/f]$ since O is valid for t . Especially, for some statement t_1 not containing method calls, if $\{s\} t_1 \{r\}$ is valid for some s and r such that $\text{vars}(s, r) \cap \{f, f', \text{assign}(t_1)\} = \emptyset$ then $\{s\} t_1[f'/f] \{r\}$ is also valid.
 - Case $f' \notin \text{fields}(t)$. The new method body is obtained by renaming f to f' . For each specification $\langle (p, q), \overline{R} \rangle \in \overline{\text{sp}}$, the modified specification $\text{renSF}(\langle (p, q), \overline{R} \rangle, f, f')$ is added to $\mathcal{T}(\text{TAE})$ if $f' \notin \text{vars}(\{p, q, \overline{R}\})$. By *IH* there is a proof outline O such that $O \vdash_{PL} t : (p, q)$ and $\overline{R} \rightarrow \text{reqs}(O)$. Since f' does not occur in O , a proof outline for $t[f'/f]$ can be constructed directly by field renaming, resulting in the specification $\text{renSF}(\langle (p, q), \overline{R} \rangle, f, f')$.
 - Case $f' \in \text{fields}(t) \wedge f \notin \text{fields}(t)$. In this case, the method body is unaltered by the rename operation. However, we discard specifications containing the renamed field f . Proof outlines for the remaining specifications follow directly by *IH*.

□

Observation B.3. Let $\text{mtds}(\overline{\text{AM}}) = \overline{\text{M}}$ and $\text{M} = \text{I } n(\overline{\text{I } \overline{x}})\{t\}$. The following properties follow from the definitions in Figure 13 and Figure 4:

1. $\text{mtds}(\text{remMtd}(\overline{\text{AM}}, m)) = \text{mtds}(\overline{\text{AM}} \setminus \overline{\text{AM}}(m)) = \overline{\text{M}} \setminus \overline{\text{M}}(m) = \text{rem}(\overline{\text{M}}, m) = \text{rem}(\text{mtds}(\overline{\text{AM}}), m)$
2. $\text{mtds}(\text{renFld}(\text{M } \overline{\text{sp}}, m, m')) = \text{mtds}(\text{I } n[m'/m](\overline{\text{I } \overline{x}})\{t\} \overline{\text{sp}}) = \text{I } n[m'/m](\overline{\text{I } \overline{x}})\{t\} = \text{ren}(\text{M}, m, m')$
3. $\text{mtds}(\text{renMtd}(\overline{\text{AM}}, m, m')) = \text{rep}(\text{mtds}(\overline{\text{AM}}), m, m')$
4. $\text{mtds}(\text{renFld}(\overline{\text{AM}}, f, f')) = \text{rep}(\text{mtds}(\overline{\text{AM}}), f, f')$

Property 3 above follows by induction over the size of $\overline{\text{AM}}$ with the base case $\text{mtds}(\text{FSrenMtd}(\emptyset, m, m')) = \emptyset = \text{rep}(\emptyset, m, m')$ and the induction step

$$\begin{aligned} & \text{mtds}(\text{renMtd}(\overline{\text{AM}} \text{ I } n(\overline{\text{I } \overline{x}})\{t\} \overline{\text{sp}}, m, m')) = \\ & \text{mtds}(\text{renMtd}(\overline{\text{AM}}, m, m')) \text{ mtds}(\text{I } n[m'/m](\overline{\text{I } \overline{x}})\{t[m'/m]\} \overline{\text{sp}}) = \\ & \text{rep}(\text{mtds}(\overline{\text{AM}}), m, m') \text{ I } n[m'/m](\overline{\text{I } \overline{x}})\{t[m'/m]\} = \text{rep}(\text{mtds}(\overline{\text{AM}} \text{ I } n(\overline{\text{I } \overline{x}})\{t\} \overline{\text{sp}}), m, m'). \end{aligned}$$

Property 4 follows by a similar argument.

Lemma B.4. Given a trait expression TAE such that $\text{TAE} \in \mathcal{T}$, then $\text{mtds}(\mathcal{T}(\text{TAE})) = \llbracket \text{TAE} \rrbracket$.

Proof. The proof is by induction over the structure of TAE .

Base case: TAE is a basic trait **trait** Tb **is** $\{\overline{F}; \overline{H}; \overline{\text{AM}}\}$ be the definition of Tb . Let $\overline{\text{M}} = \text{mtds}(\overline{\text{AM}})$. This trait is introduced in $\mathcal{T}$ by rule **BASICTRAIT**. By induction over the inference rules, we then know that $\mathcal{T}(\text{Tb}) = \{\overline{F}; \overline{H}; \overline{\text{AM}}'\}$ for some $\overline{\text{AM}}'$ such that $\text{mtds}(\overline{\text{AM}}') = \text{mtds}(\overline{\text{AM}})$, since no rules add new method definitions to a trait in $\mathcal{T}$. (However, $\overline{\text{AM}}'$ may contain more specifications than $\overline{\text{AM}}$.) From the definition of flattening for a basic trait (Figure 3), we get

$$\text{mtds}(\mathcal{T}(\text{Tb})) = \text{mtds}(\overline{\text{AM}}') = \overline{\text{M}} \stackrel{\text{Figure 3}}{=} \llbracket \{\overline{F}; \overline{H}; \overline{\text{M}}\} \rrbracket \stackrel{\text{Figure 3}}{=} \llbracket \text{Tb} \rrbracket.$$

Induction step: $\text{TAE} = \text{TAE}' \text{ ao}$. Since $\text{TAE} \in \mathcal{T}$, we know from Lemma B.1 that $\text{TAE}' \in \mathcal{T}$. Assume that $\mathcal{T}(\text{TAE}') = \{\bar{F}; \bar{H}; \bar{AM}\}$, so $\text{mtds}(\mathcal{T}(\text{TAE}')) = \text{mtds}(\bar{AM})$. The induction hypothesis *IH* is that $\text{mtds}(\mathcal{T}(\text{TAE}')) = \llbracket \text{TAE}' \rrbracket$, so, $\llbracket \text{TAE}' \rrbracket = \text{mtds}(\bar{AM})$. The different trait alteration operations are treated as separate cases:

- $\text{ao} = [\text{exclude } m]$:

$$\begin{aligned} \text{mtds}(\mathcal{T}(\text{TAE})) &\stackrel{\text{Figure 13}}{=} \text{mtds}(\text{remMtd}(\bar{AM}, m)) \stackrel{\text{Obs. B.3}}{=} \text{rem}(\text{mtds}(\bar{AM}), m) \stackrel{\text{IH}}{=} \\ &\text{rem}(\llbracket \text{TAE}' \rrbracket, m) \stackrel{\text{Figure 3}}{=} \llbracket \text{TAE}'[\text{exclude } m] \rrbracket = \llbracket \text{TAE} \rrbracket \end{aligned}$$

- $\text{ao} = [m \text{ aliasAs } m']$. Case $m \notin \bar{AM}$:

$$\text{mtds}(\mathcal{T}(\text{TAE})) \stackrel{\text{Figure 13}}{=} \text{mtds}(\bar{AM}) \stackrel{\text{IH}}{=} \llbracket \text{TAE}' \rrbracket \stackrel{\text{Figure 3}}{=} \llbracket \text{TAE}'[m \text{ aliasAs } m'] \rrbracket = \llbracket \text{TAE} \rrbracket$$

Case $m \in \bar{AM}$:

$$\begin{aligned} \text{mtds}(\mathcal{T}(\text{TAE})) &\stackrel{\text{Figure 13}}{=} \text{mtds}(\bar{AM}) \text{ mtds}(\text{renId}(\bar{AM}(m), m, m')) \stackrel{\text{Obs. B.3}}{=} \text{mtds}(\bar{AM}) \text{ ren}(\bar{AM}(m), m, m') \\ &\stackrel{\text{IH}}{=} \llbracket \text{TAE}' \rrbracket \text{ ren}(\llbracket \text{TAE}' \rrbracket(m), m, m') \stackrel{\text{Figure 3}}{=} \llbracket \text{TAE}[m \text{ aliasAs } m'] \rrbracket = \llbracket \text{TAE} \rrbracket \end{aligned}$$

- $\text{ao} = [m \text{ renameTo } m']$:

$$\begin{aligned} \text{mtds}(\mathcal{T}(\text{TAE})) &\stackrel{\text{Figure 13}}{=} \text{mtds}(\text{renMtd}(\bar{AM}, m, m')) \stackrel{\text{Obs. B.3}}{=} \text{rep}(\text{mtds}(\bar{AM}), m, m') \stackrel{\text{IH}}{=} \text{rep}(\llbracket \text{TAE}' \rrbracket, m, m') \\ &\stackrel{\text{Figure 3}}{=} \llbracket \text{TAE}'[m \text{ renameTo } m'] \rrbracket = \llbracket \text{TAE} \rrbracket \end{aligned}$$

- $\text{ao} = [f \text{ renameTo } f']$:

$$\begin{aligned} \text{mtds}(\mathcal{T}(\text{TAE})) &\stackrel{\text{Figure 13}}{=} \text{mtds}(\text{renFld}(\bar{AM}, f, f')) \stackrel{\text{Obs. B.3}}{=} \text{rep}(\text{mtds}(\bar{AM}), f, f') \stackrel{\text{IH}}{=} \text{rep}(\llbracket \text{TAE}' \rrbracket, f, f') \\ &\stackrel{\text{Figure 3}}{=} \llbracket \text{TAE}'[f \text{ renameTo } f'] \rrbracket = \llbracket \text{TAE} \rrbracket \end{aligned}$$

□

Lemma B.5. Let CTE be a composed trait expression defined by $\text{TAE}_1 + \dots + \text{TAE}_n$, and assume that each $\text{TAE}_i \in \mathcal{T}$. Then $\text{mtds}(\text{CTE}) = \llbracket \text{CTE} \rrbracket$.

Proof. This follows directly from the definition of $\text{mtds}(\text{CTE})$ together with Lemma B.4 and the definition of flattening (Section 3.2).

$$\text{mtds}(\text{CTE}) \triangleq \bigcup_{1 \leq i \leq n} \text{mtds}(\mathcal{T}(\text{TAE}_i)) \stackrel{\text{Lemma B.4}}{=} \bigcup_{1 \leq i \leq n} \llbracket \text{TAE}_i \rrbracket \stackrel{\text{Figure 3}}{=} \llbracket \text{CTE} \rrbracket$$

□

Lemma B.6. Consider **class** C **implements** $\bar{I}$ **by** $\{\bar{rp}_I \bar{F};\}$ **and** CTE such that $\text{CTE} \triangleq \text{TAE}_1 + \dots + \text{TAE}_n$. Let the environments $\mathcal{T}$ and $\mathcal{C}$ result from the successful analysis of this class. For $m \in \text{mtds}(\text{TAE}_i)$ and $1 \leq i \leq n$, we have

$$S_{\mathcal{C}}(C, m) \subseteq \text{specs}(\mathcal{T}(\text{TAE}_i)(m)).$$

Proof. By type safety, each method in CTE is defined in exactly one TAE_i . The relation $S_{\mathcal{C}}(C, m) \subseteq \text{specs}(\mathcal{T}(\text{TAE}_i)(m))$ follows by induction over the inference rules of PST(PL). When the analysis of class C starts by rule CLASS, the mapping $S_{\mathcal{C}}(C, m)$ is initially empty, so the subset relation holds initially. Since $S_{\mathcal{C}}(C, m)$ is only extended by ANALYZE, it suffices to consider this rule. The rule maintains the relation since if $S_{\mathcal{C}}(C, m)$ is extended by some set $\bar{sp}$ by the rule, it follows from the premises that $\bar{sp} \subseteq \text{specs}(\mathcal{T}(\text{TAE}_i)(m))$. □

C. Verification of the Example

This section details the verification of the bank account classes from Section 8. In order to verify proof outlines, we must fix the program logic PL . In this example, we use a standard Hoare logic for sequential statements [AdBO09].

Since formal parameters are read-only in the example, we may use the following rule for analyzing method calls:

$$\frac{\text{(HOARECALL)} \quad \{p\} \text{ m } (\bar{x}) \quad \{q\}}{\{p[\bar{e}/\bar{x}]\} \text{ v=m } (\bar{e}) \quad \{q[\bar{e}/\bar{x}][v/\text{result}]\}}$$

where $\bar{x}$ are the formal parameters of the method m . Note that the annotations of a call $\text{v=m } (\bar{e})$ may be strengthened by assertions over logical and method local variables. Let $\mathcal{X}$ and $\mathcal{Z}$ denote the sets of method local variables (including formal parameters) and logical variables, respectively. If $m : (p, q) \in \text{reqs}(O)$ is a requirement in some proof outline O , then O may contain decorated call statements of the form

$$\{p[\bar{e}/\bar{x}] \wedge c\} \text{ v=m } (\bar{e}) \quad \{q[\bar{e}/\bar{x}][v/\text{result}] \wedge c\}$$

where $\text{vars}(c) \subseteq \{\mathcal{X} \cup \mathcal{Z}\}$.

Initially in an analysis, the trait and class environments are empty. Let $\mathcal{T}_0$ be the trait environment resulting from the analysis of all the basic traits in Figure 8, so each basic trait is bound to an annotated basic trait expression in $\mathcal{T}_0$. These expressions correspond to the definitions in Figure 8, for example:

$$\mathcal{T}_0(\text{TBasicUpd}) = \{\text{int } \text{bal}; \emptyset; \text{void } \text{update}(\text{nat } x)\{ \text{bal} = \text{bal} - x \} \langle (bal = b_0, bal = b_0 - x), \emptyset \rangle\}.$$

The only proof outlines which need to be supplied during the basic trait analysis are for this `update` method and for `deposit` in `TBasicAccount`. It is straightforward to provide these proof outlines given the guarantees.

We now consider the analysis of the class `CFeeAccount` in detail, assuming that the class is analyzed in the environment $\mathcal{C}_0, \mathcal{T}_0$, where `CFeeAccount` $\notin \mathcal{C}_0$. Figure 14 shows the application of the proof system to the analysis of this class, leading to a successful analysis, where $\mathcal{C}_5, \mathcal{T}_3$ are the resulting environments.

In the first step, the class definition leads to the analysis of the trait expression `TBasicAccount + CTE` followed by `discharge` operations for the contracts of the interface `IFeeAccount`. The trait expression decomposes and as the basic traits are already analyzed, we get to the `extend` operation for the trait alteration expression `TBasicUpd[update rt bUpdate]`. The result is an extended trait environment $\mathcal{T}_1$ where a trait expression for the renaming of `TBasicUpd` is included in the trait environment. It is not necessary to reanalyze the renamed method, the specification is generated by renaming the existing method specification in `TBasicUpd`.

At this point, the `discharge` operation is at the head of the analysis operation sequence, and is decomposed. For the method `deposit`, `ANALYZE` can be applied and the specification found in `TBasicAccount` is included in the class environment, resulting in the class environment $\mathcal{C}_2$. Since neither `withdraw` nor `update` is specified in `TBasicAccount` and `TFeeUpd`, respectively, the analysis continues with `OPENANALYSIS`. Proof outlines for these methods must be provided at the level of their trait declarations by applying `INCREMENT` and checked by applying `VERIFY`. In these proof outlines, the requirement for the call to `bUpdate` follows from the generated specification of this method.

$$\begin{array}{c}
\frac{\mathcal{C}_5, \mathcal{T}_3 \vdash}{\mathcal{C}_5, \mathcal{T}_3 \vdash \langle \text{CFeeAccount} : \varepsilon \rangle} \text{(EMPCCLASS)} \\
\frac{\mathcal{C}_5, \mathcal{T}_3 \vdash \langle \text{CFeeAccount} : \text{discharge}(\mathbf{0}) \rangle}{\mathcal{C}_5, \mathcal{T}_3 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}], \text{bUpdate} : (bal == b_0, bal == b_0 - x)) \rangle} \text{(EMPDISCHARGE)} \\
\frac{\mathcal{C}_4, \mathcal{T}_3 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}], \text{bUpdate} : (bal == b_0, bal == b_0 - x)) \rangle}{\mathcal{C}_4, \mathcal{T}_3 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{bUpdate} : (bal == b_0, bal == b_0 - x)) \rangle} \text{(ANALYZE) (14)} \\
\frac{\mathcal{C}_4, \mathcal{T}_3 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{bUpdate} : (bal == b_0, bal == b_0 - x)) \rangle}{\mathcal{C}_3, \mathcal{T}_3 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TFeeUpd}, \text{update} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(OPENANALYSIS) (13)} \\
\frac{\mathcal{C}_3, \mathcal{T}_3 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TFeeUpd}, \text{update} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_3, \mathcal{T}_2 \vdash \text{verify}(\text{TFeeUpd}, \text{U} \langle (bal == b_0, bal == b_0 - x - fee), \text{bUpdate} : (bal == b_0, bal == b_0 - x) \rangle) \cdot \langle \text{CFeeAccount} : \dots \rangle} \text{(VERIFY) (11)} \\
\frac{\mathcal{C}_3, \mathcal{T}_2 \vdash \text{verify}(\text{TFeeUpd}, \text{U} \langle (bal == b_0, bal == b_0 - x - fee), \text{bUpdate} : (bal == b_0, bal == b_0 - x) \rangle) \cdot \langle \text{CFeeAccount} : \dots \rangle}{\mathcal{C}_3, \mathcal{T}_2 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TFeeUpd}, \text{update} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(INCREMENT) (10)} \\
\frac{\mathcal{C}_3, \mathcal{T}_2 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TFeeUpd}, \text{update} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_3, \mathcal{T}_2 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{update} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(OPENANALYSIS) (9)} \\
\frac{\mathcal{C}_3, \mathcal{T}_2 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{update} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_2, \mathcal{T}_2 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TBasicAccount}, \text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(ANALYZE) (8)} \\
\frac{\mathcal{C}_2, \mathcal{T}_2 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TBasicAccount}, \text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_2, \mathcal{T}_1 \vdash \text{verify}(\text{TBasicAccount}, \text{W} \langle (bal == b_0, bal == b_0 - x - fee), \text{update} : (bal == b_0, bal == b_0 - x - fee) \rangle) \cdot \langle \text{CFeeAccount} : \dots \rangle} \text{(VERIFY) (7)} \\
\frac{\mathcal{C}_2, \mathcal{T}_1 \vdash \text{verify}(\text{TBasicAccount}, \text{W} \langle (bal == b_0, bal == b_0 - x - fee), \text{update} : (bal == b_0, bal == b_0 - x - fee) \rangle) \cdot \langle \text{CFeeAccount} : \dots \rangle}{\mathcal{C}_2, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TBasicAccount}, \text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(INCREMENT) (6)} \\
\frac{\mathcal{C}_2, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TBasicAccount}, \text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_2, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(OPENANALYSIS) (5)} \\
\frac{\mathcal{C}_2, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_2, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{discharge}(\mathbf{0}) \cdot \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(EMPDISCHARGE)} \\
\frac{\mathcal{C}_2, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{discharge}(\mathbf{0}) \cdot \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_1, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TBasicAccount}, \text{deposit} : (bal == b_0, bal == b_0 + x)) \cdot \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(ANALYZE) (4)} \\
\frac{\mathcal{C}_1, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{analyze}(\text{TBasicAccount}, \text{deposit} : (bal == b_0, bal == b_0 + x)) \cdot \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_1, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{deposit} : (bal == b_0, bal == b_0 + x)) \cdot \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(OPENANALYSIS) (3)} \\
\frac{\mathcal{C}_1, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{deposit} : (bal == b_0, bal == b_0 + x)) \cdot \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_1, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{deposit} : (bal == b_0, bal == b_0 + x)) \text{U} \text{withdraw} : (bal == b_0, bal == b_0 - x - fee) \rangle} \text{(DECOMP4)} \\
\frac{\mathcal{C}_1, \mathcal{T}_1 \vdash \langle \text{CFeeAccount} : \text{discharge}(\text{deposit} : (bal == b_0, bal == b_0 + x)) \text{U} \text{withdraw} : (bal == b_0, bal == b_0 - x - fee) \rangle}{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}], \langle \text{CFeeAccount} : \dots \rangle)} \text{(EXTEND) (2)} \\
\frac{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}], \langle \text{CFeeAccount} : \dots \rangle)}{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TFeeUpd}, \text{extend}(\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}], \langle \text{CFeeAccount} : \dots \rangle))} \text{(LOOKUP)} \\
\frac{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TFeeUpd}, \text{extend}(\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}], \langle \text{CFeeAccount} : \dots \rangle))}{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TBasicUpd} + \text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}], \langle \text{CFeeAccount} : \dots \rangle)} \text{(DECOMP3)} \\
\frac{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TBasicUpd} + \text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}], \langle \text{CFeeAccount} : \dots \rangle)}{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TBasicAccount}) \cdot \text{extend}(\text{CTE}) \cdot \langle \text{CFeeAccount} : \dots \rangle} \text{(LOOKUP)} \\
\frac{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TBasicAccount}) \cdot \text{extend}(\text{CTE}) \cdot \langle \text{CFeeAccount} : \dots \rangle}{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TBasicAccount} + \text{CTE}) \cdot \langle \text{CFeeAccount} : \text{discharge}(\text{deposit} : (bal == b_0, bal == b_0 + x)) \cdot \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle} \text{(DECOMP3)} \\
\frac{\mathcal{C}_1, \mathcal{T}_0 \vdash \text{extend}(\text{TBasicAccount} + \text{CTE}) \cdot \langle \text{CFeeAccount} : \text{discharge}(\text{deposit} : (bal == b_0, bal == b_0 + x)) \cdot \text{discharge}(\text{withdraw} : (bal == b_0, bal == b_0 - x - fee)) \rangle}{\mathcal{C}_0, \mathcal{T}_0 \vdash \text{class CFeeAccount implements IFeeAccount by \{int bal, nat fee\} and TBasicAccount + CTE}} \text{(CLASS) (1)}
\end{array}$$

Side conditions and auxiliary computations in the proof:

- (1) $\mathcal{C}_1 = \langle D_{\mathcal{C}_0}[\text{CFeeAccount} \mapsto \langle \text{IFeeAccount}, \text{TBasicAccount} + \text{CTE}, \text{int bal}, \text{nat fee} \rangle], S_{\mathcal{C}_0} \rangle$
- (2) $\mathcal{T}_1 = \mathcal{T}_0[\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}] \mapsto \text{mod}(\mathcal{T}_0(\text{TBasicUpd}), [\text{update } \mathbf{rT} \text{ bUpdate}])]$
 where $\text{mod}(\mathcal{T}_0(\text{TBasicUpd}), [\text{update } \mathbf{rT} \text{ bUpdate}])$ evaluates to:
 $\{\text{int bal}; \mathbf{0}; \text{void bUpdate}(\text{nat x})\{bal == bal - x\} \langle (bal == b_0, bal == b_0 - x), \mathbf{0} \rangle\}$
- (3) $\text{TBasicAccount} \in \text{impl}(D_{\mathcal{C}_1}(\text{CFeeAccount})) \wedge \text{deposit} \in \mathcal{T}_1(\text{TBasicAccount})$
- (4) $\langle (bal == b_0, bal == b_0 + x), \mathbf{0} \rangle \in \text{specs}(\mathcal{T}_1(\text{TBasicAccount})(\text{deposit}))$
 $\mathcal{C}_2 = \langle D_{\mathcal{C}_1}, S_{\mathcal{C}_1}[(\text{CFeeAccount}, \text{deposit}) \mapsto \langle (bal == b_0, bal == b_0 + x), \mathbf{0} \rangle] \rangle$
- (5) $\text{TBasicAccount} \in \text{impl}(D_{\mathcal{C}_2}(\text{CFeeAccount})) \wedge \text{withdraw} \in \mathcal{T}_1(\text{TBasicAccount})$
- (6) $\text{W} = \text{void withdraw}(\text{nat x})\{\text{update}(x)\}$
- (7) $\mathcal{T}_2 = \mathcal{T}_1[\text{TBasicAccount} \mapsto \text{addSpec}(\mathcal{T}_1(\text{TBasicAccount}), \text{withdraw}, \langle (bal == b_0, bal == b_0 - x - fee), \text{update} : (bal == b_0, bal == b_0 - x - fee) \rangle)]$
- (8) $\langle (bal == b_0, bal == b_0 - x - fee), \text{update} : (bal == b_0, bal == b_0 - x - fee) \rangle \in \text{specs}(\mathcal{T}_2(\text{TBasicAccount})(\text{withdraw}))$
 $\mathcal{C}_3 = \langle D_{\mathcal{C}_2}, S_{\mathcal{C}_2}[(\text{CFeeAccount}, \text{withdraw}) \mapsto \langle (bal == b_0, bal == b_0 - x - fee), \text{update} : (bal == b_0, bal == b_0 - x - fee) \rangle] \rangle$
- (9) $\text{TFeeUpd} \in \text{impl}(D_{\mathcal{C}_3}(\text{CFeeAccount})) \wedge \text{update} \in \mathcal{T}_2(\text{TFeeUpd})$
- (10) $\text{U} = \text{void update}(\text{nat x})\{\text{bUpdate}(x + fee)\}$
- (11) $\mathcal{T}_3 = \mathcal{T}_2[\text{TFeeUpd} \mapsto \text{addSpec}(\mathcal{T}_2(\text{TFeeUpd}), \text{update}, \langle (bal == b_0, bal == b_0 - x - fee), \text{bUpdate} : (bal == b_0, bal == b_0 - x) \rangle)]$
 and the proof outline follows directly by substituting formal parameters in the requirement according to rule HOARECALL
- (12) $\langle (bal == b_0, bal == b_0 - x - fee), \text{bUpdate} : (bal == b_0, bal == b_0 - x) \rangle \in \text{specs}(\mathcal{T}_3(\text{TFeeUpd})(\text{update}))$
 $\mathcal{C}_4 = \langle D_{\mathcal{C}_3}, S_{\mathcal{C}_3}[(\text{CFeeAccount}, \text{update}) \mapsto \langle (bal == b_0, bal == b_0 - x - fee), \text{bUpdate} : (bal == b_0, bal == b_0 - x) \rangle] \rangle$
- (13) $\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}] \in \text{impl}(D_{\mathcal{C}_4}(\text{CFeeAccount})) \wedge \text{bUpdate} \in \mathcal{T}_3(\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}])$
- (14) $\langle (bal == b_0, bal == b_0 - x), \mathbf{0} \rangle \in \text{specs}(\mathcal{T}_3(\text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}])(\text{update}))$
 $\mathcal{C}_5 = \langle D_{\mathcal{C}_4}, S_{\mathcal{C}_4}[(\text{CFeeAccount}, \text{bUpdate}) \mapsto \langle (bal == b_0, bal == b_0 - x), \mathbf{0} \rangle] \rangle$

Fig. 14. Analysis details for the class CFeeAccount. To simplify the presentation, side conditions and auxiliary computations for the different rule applications are given as notes, CTE abbreviates the expression $\text{TFeeUpd} + \text{TBasicUpd}[\text{update } \mathbf{rT} \text{ bUpdate}]$, $\mathbf{rT}$ abbreviates renameTo , and $\langle \text{CFeeAccount} : \dots \rangle$ denotes that the class contains the same operations as found on the line below.