

Traits: Correctness-by-Construction for Free

Tobias Runge^{1,2}, Alex Potanin³, Thomas Thüm⁴, and Ina Schaefer^{1,2}

¹ TU Braunschweig, Germany

² Karlsruhe Institute of Technology, Germany

³ Australian National University, Australia

⁴ University of Ulm, Germany

{tobias.runge,ina.schaefer}@kit.edu, alex.potanin@anu.edu.au,
thomas.thuem@uni-ulm.de

Abstract. We demonstrate that traits are a natural way to support correctness-by-construction (CbC) in an existing programming language in the presence of traditional post-hoc verification (PhV). With Correctness-by-Construction, programs are constructed incrementally along with a specification that is inherently guaranteed to be satisfied. CbC is complex to use without specialized tool support, since it needs a set of refinement rules of fixed granularity which are additional rules on top of the programming language.

In this work, we propose TraitCbC, an incremental program construction procedure that implements correctness-by-construction on the basis of PhV by using traits. TraitCbC enables program construction by trait composition instead of refinement rules. It provides a programming guideline, which similar to CbC should lead to well-structured programs, and allows flexible reuse of verified program building blocks. We introduce TraitCbC formally and prove the soundness of our verification strategy. Additionally, we implement TraitCbC as a proof of concept.

1 Introduction

Correctness-by-Construction (CbC) [19, 22, 30, 37] is a methodology that incrementally constructs correct programs guided by a pre-/postcondition specification.⁵ CbC uses small tractable refinement rules where in each refinement step, an abstract statement (i.e., a hole in the program) is refined to a more concrete implementation that can still contain some nested abstract statements. While refining the program, the correctness of the whole program is guaranteed through the check of conditions in the refinement rules. The construction ends when no abstract statement is left. Through the structured reasoning discipline that is

⁵ The approach should not be confused with other CbC approaches such as CbyC of Hall and Chapman [24]. CbyC is a software development process that uses formal modeling techniques and analysis for various stages of development (architectural design, detailed design, code) to detect and eliminate defects as early as possible [13]. We also exclude data refinement from abstract data types to concrete ones during code generation as for example in Isabelle/HOL [23].

enforced by the refinement rules, it is claimed that program quality increases and verification effort is reduced [30, 50].

Despite these benefits, CbC has a drawback: the refinement rules extend the programming language (i.e., refinements are an additional linguistic construct to transform programs). Special tool support [42] is necessary to introduce the CbC refinement process to a programming language. Additionally, the predefined rules have a fine granularity such that for every new statement the programmer adds to the program, an application of a refinement rule is necessary. Consequently, the concepts of CbC (e.g., abstract statements and refinement rules) increase the effort and necessary knowledge of the developer to construct programs.

Post-hoc verification (PhV) is another approach to develop correct programs. A method is verified against its pre- and postconditions after implementation. In practice, it often happens that a program is constructed first, with the objective of verifying it later [50]. This can lead to tedious verification work if the program is not well-structured. An example is the difficult search for the many reasons preventing the verification of a method to be completed: an incorrect specification, an incorrect method, or inadequate tool support. Therefore, a structured programming approach is desirable to construct programs which are amenable to software verification.

In this work, we use *traits* [20] to overcome the drawbacks of CbC (complex programming style using external refinement rules) and introduce a programming guideline for an incremental trait-based program construction approach that guarantees that the resulting trait-based program is correct-by-construction. TraitCbC is based on PhV. With TraitCbC, the same programs can be verified as with PhV, but in addition, TraitCbC introduces an explicit program construction approach. It utilizes the flexibility of traits, which is beneficial for scenarios as incremental development [18] and the development of software product lines [15, 10].

Traits [20] are a flexible object-oriented language construct supporting a rich form of modular code reuse orthogonal to inheritance. A trait contains a set of *concrete* or *abstract* methods (i.e., the method has either a body or has no body), independent of any class or inheritance hierarchy.⁶ Traits are independent modules that can be composed into larger traits or classes. When traits are composed, the resulting code contains all methods of all composed traits. To verify traits, Damiani et al. [18] proposed a modular and incremental *post-hoc verification* process. Each method in every trait is verified in isolation by showing that the method satisfies its *contract* [35]. Then, during the composition of traits, it has to be checked whether a method implemented in one trait is compatible with the abstract method with the same signature in another trait. That means, a concrete method has to satisfy the specification of the abstract method. A concrete method with a weaker precondition and a stronger postcondition fulfills the contract of the abstract method (cf. Liskov substitution principle [34]).

⁶ The term trait has been used by many programming languages: Java interfaces with default methods are a good approximation for what has been called trait in the literature, while Scala traits are mixins [21], and Rust traits are type classes [46].

A developer using TraitCbC starts by implementing a method (e.g., a method **a**) in a first trait. Similar to CbC, the method can contain holes that are refined in subsequent steps. A hole in TraitCbC is an abstract method (e.g., an abstract method **b**) that is called in method **a**; that is, a call to an abstract method corresponds to an abstract statement in CbC. In the next step, one of these new abstract methods (e.g., **b**) is implemented in a second trait, again more abstract methods can be declared for the implementation. Similar to PhV, it must be proven that the implemented methods satisfy their specifications. Afterwards, the traits are composed; the composition operation checks that the contract of the concrete method **b** in the second trait fulfills the contract of the abstract method **b** in the first trait. This incremental process stops when the last abstract method is implemented, and all traits are composed.

The main result of our work is the discovery that traits intrinsically enable correctness-by-construction. This work is not about pushing verification forward in the sense of adding more expressive power. TraitCbC realizes a refinement-based program development approach using pre-/postcondition contracts and method calls instead of refinement rules and abstract statements as in CbC. Refinement rules in the form of trait composition exist as a direct concept of the programming language instead of being a program transformation concept. Additionally, each method implemented in the refinement process can be reused by composing traits in different contexts (i.e., already proven methods can be called by new methods under construction). This is advantageous compared to the limited reuse potential of methods in class-based inheritance. Finally, TraitCbC is parametric w.r.t. the specification logic. Thus, a language with traits can adopt the proposed CbC methodology.

2 Motivating Example

In this section, we go through an example of how our development process enables CbC using traits.

Incremental Construction of MaxElement We use a sample object-oriented language in the code examples. We construct a method `maxElement` that finds the maximum element in a list of numbers. A list has a head and a tail. Only non-empty lists have a maximum element. This is explicit in the precondition of our specification, where we require that the list has at least one element. In the postcondition, we specify that the result is in the list and larger than or equal to every other element. A method `contains` checks that the result is a member of the list. In the first step, we create a trait `MaxETrait1` that defines the abstract method `maxElement`. The method `maxElement` is abstract, i.e., equivalent to an abstract statement in CbC.

```

1 trait MaxETrait1 {
2   @Pre: list.size() > 0
3   @Post: list.contains(result) &
4     (forall Num n: list.contains(n) ==> result >= n)

```

```

5   abstract Num maxElement(List list);
6 }

```

In the second step in trait `MaxETrait2`, we implement the method `maxElement` using two abstract methods. We introduce an `if-elseif-else`-expression where the branches invoke abstract methods. The guards check whether the list has only one element or whether the current element is larger than or equal to the maximum of the rest of the list. The abstract method `accessHead` returns the current element, and the abstract method `maxTail` returns the maximum in the remaining list. So, we recursively search the list for the largest element by comparing the maximum element of the list tail with the current element until we reach the end of the list.

```

1 trait MaxETrait2 {
2   @Pre: list.size() > 0
3   @Post: list.contains(result) &
4     (forall Num n: list.contains(n) ==> result >= n)
5   Num maxElement(List list) =
6     if (list.size() == 1) {accessHead(list)}
7     elseif (accessHead(list) >= maxTail(list))
8       {accessHead(list)}
9     else {maxTail(list)}
10
11  @Pre: list.size() > 0
12  @Post: result == list.element()
13  abstract Num accessHead(List list);
14
15  @Pre: list.size() > 1
16  @Post: list.tail().contains(result) &
17    (forall Num n: list.tail().contains(n) ==> result >= n)
18  abstract Num maxTail(List list);
19 }

```

The correct implementation of the method `maxElement` can be guaranteed under the assumptions that all introduced abstract methods are correctly implemented. Similar to PhV, a program verifier conducts a proof of method `maxElement` and uses the introduced specifications of the methods `accessHead` and `maxTail`. When the proof succeeds, we know that the first method is correctly implemented. In our incremental CbCTrait process, we verify each method implementation directly after construction; and so we are able to reuse each implemented method in the following steps (e.g., by calling the method in the body of other methods).

We now compose the developed traits to complete the first refinement step. To perform the composition `MaxETrait1 + MaxETrait2`, we check that the specification of the method `maxElement` fulfills the specification of the abstract method in the first trait (cf. Liskov substitution principle [34]). In this case, this means checking that:

```

MaxETrait1.maxElement(..).pre ==> MaxETrait2.maxElement(..).pre as well as:
MaxETrait2.maxElement(..).post ==> MaxETrait1.maxElement(..).post.

```

When the composition of two verified traits is successful, the result is also a verified trait. Note that the composed trait does not need to be verified directly by a program verifier in TraitCbC because it is correct by construction. In this example, the specifications are the same, thus checking for a successful composition is trivial, but this is not generally the case. In particular, the logic needs to take into account ill-founded specifications and recursion in the specification. We discuss the difficulties of handling those cases in the technical report [41].

The methods `accessHead` and `maxTail` are implemented in the next two refinement steps in traits `MaxETrait3` and `MaxETrait4`⁷. As we implement a recursive method, the method `maxTail` calls the `maxElement` method, thus `maxElement` is introduced as an abstract method in this trait. We have to verify that the method `accessHead` satisfies its specification using a program verifier. Similarly, we have to verify the correctness of the method `maxTail`.

```

1 trait MaxETrait3 {
2   @Pre: list.size() > 0
3   @Post: result == list.element()
4   Num accessHead(List list) = list.element()
5 }

1 trait MaxETrait4 {
2   @Pre: list.size() > 1
3   @Post: list.tail().contains(result) &
4     (forall Num n: list.tail().contains(n) ==> result >= n)
5   Num maxTail(List list) = maxElement(list.tail())
6
7   @Pre: list.size() > 0
8   @Post: list.contains(result) &
9     (forall Num n: list.contains(n) ==> result >= n)
10  abstract Num maxElement(List list);
11 }

```

As before, all traits are composed, and it is checked that the specifications of the concrete methods fulfill the specifications of the abstract ones. As we have no contradicting specifications for the same methods, the composition is well-formed. The final program `MaxE` is as follows.

```
1 class MaxE = MaxETrait1 + MaxETrait2 + MaxETrait3 + MaxETrait4
```

Advantages of TraitCbC As shown in the example, TraitCbC enables the CbC programming style without the need of external refinement rules. In classical CbC, when designing a unit of code, the programmer has to proceed with atomic steps of a predefined granularity. In contrast, in TraitCbC the programmer is free to divide a unit of code in any granularity, by including as many auxiliary methods as needed to bring the verification to an appropriate granularity. TraitCbC helps to construct code in fine-grained steps which are more amenable for verification than single more complex methods. If the programmer chooses

⁷ The methods could also be implemented in one trait.

to not include any auxiliary methods at all, this is essentially the same as the traditional post-hoc verification style. In the example above, we could implement the method `maxElement` in one step without the intermediate step that introduces the two abstract methods `accessHead` and `maxTail`.

Additionally, the already proven auxiliary methods in traits can be reused. For example, if we want to implement a `minElement` method, we could reuse already implemented traits to reduce the programming and verification effort. The method `minElement` is implemented in the following in trait `MinE` with one abstract method. The specification of the method `accessHead` is the same as for the method `accessHead` above, so `MaxETrait3` can be reused. In this example, we show the flexible granularity of `TraitCbC` by directly implementing the else branch, instead of introducing an auxiliary method as for `maxElement`.

```

1  trait MinE {
2    @Pre: list.size() > 0
3    @Post: list.contains(result) &
4      (forall Num n: list.contains(n) ==> result <= n)
5    Num minElement(List list) =
6      if (list.size() == 1) {accessHead(list)}
7      elseif (accessHead(list) <= minElement(list.tail()))
8        {accessHead(list)}
9      else {minElement(list.tail())}
10
11   @Pre: list.size() > 0
12   @Post: result == list.element()
13   abstract Num accessHead(List list);
14 }

```

The correctness of `minElement` is verified with the specifications of the method `accessHead`. By composing `MinE` with `MaxETrait3`, we get a correct implementation of `minElement`. Note how this verification process supports abstraction: as long as the contracts are compatible, methods can be implemented in different styles by different programmers to best meet non-functional requirements while preserving the specified observable behavior [9]. A completely different implementation of `maxElement` can be used if it fulfills the specification of the abstract method `maxElement` in trait `MaxETrait1`. This decoupling of specification and corresponding satisfying implementations facilitates an incremental development process where a specified code base is extended with suitable implementations [18].

3 Object-Oriented Trait-Based Language

In this section, we formally introduce the syntax, type system, and flattening semantics of a minimal core calculus for `TraitCbC`. We keep this calculus for `TraitCbC` parametric in the specification logic so that it can be used with a suitable program verifier and associated logic. The presented rules to compose traits are conventional. The focus of our work is to enable a `CbC` approach using traits that programmers can easily adopt. Therefore, we present the calculus to

prove soundness of TraitCbC, but focus on the presentation of the advantages of incremental trait-based programming in this paper. Indeed, languages with traits and with a suitable specification language intrinsically enable incremental program construction. For the sake of completeness, reduction rules of TraitCbC are presented in the technical report [41].

3.1 Syntax

The concrete syntax of our core calculus for TraitCbC is shown in Fig. 1, where non-terminals ending with ‘s’ are implicitly defined as a sequence of non-terminals, i.e., $vs ::= v_1 \dots v_n$. We use the metavariables t for trait names, C for class names and m for method names. A program consists of trait and class definitions. Each definition has a name and a trait expression E . The trait expression can be a *Body*, a trait name, a composition of two trait expressions E , or a trait expression E where a method is made abstract, written as $E[\text{makeAbstract } m]$. A *Body* has a flag **interface** to define an interface, a set of implemented interfaces Cs and a list of methods Ms . Methods have a method header MH consisting of a specification S , the return type, a method name, and a list of parameters. Methods have an optional method body. In the method body, we have standard expressions, such as variable references, method calls, and object initializations. For simplicity, we exclude updatable state. Field declarations are emulated by method declarations, and field accesses are emulated by method calls.

The specification S in each method header is used to verify that methods are correctly implemented. The specification is written in some logic. In our examples, we will use first-order logic (cf. the example in Section 2). A well-formed program respects the following conditions:

Every *Name* in Ds must be unique so that Ds can be seen as a map from names to trait expressions. Trait expressions E can refer to trait names t . A well-formed Ds does not have any circular trait definitions like $t = t$ or $t_1 = t_2$ and $t_2 = t_1$. In a *Body*, all names of implemented interfaces must be unique and all method names must be unique, so that *Body* is a map from method names to method definitions. In a method header, parameters must have unique names, and no explicit parameter can be called **this**.

3.2 Typing Rules

In our type system, we have a typing context $\Gamma ::= x_1 : C_1 \dots x_n : C_n$ which assigns types C_i to variables x_i . We define typing rules for our three kinds of expressions: x , method calls, and object initialization. We combine typing and verification in our type checking $\Gamma \vdash e : C \dashv P_0 \models P_1$. This judgment can be read as: under typing context Γ , the expression e has type C , where under the knowledge P_0 we need to prove P_1 . The knowledge P_0 is our collected information that we use to prove a method correct. That means, in our typing rules, we collect the knowledge about the parameters and expressions in a method body to verify that this method body fulfills the specification defined in the method header. The verification obligation P_1 should follow from the knowledge P_0 .

$Prog$	$::= Ds e$
D	$::= TD \mid CD$
$Name$	$::= t \mid C$
TD	$::= t = E$
CD	$::= C = E$
E	$::= Body \mid t \mid E + E \mid E[\mathbf{makeAbstract} \ m]$
$Body$	$::= \{\mathbf{interface?} \ [Cs] \ Ms\}$
M	$::= MH \ e?;$
MH	$::= S \ \mathbf{method} \ C \ m(C_1 \ x_1 \dots C_n \ x_n)$
e	$::= x \mid e.m(es) \mid \mathbf{new} \ C(es)$
\mathcal{E}_v	$::= [] .m(es) \mid v.m(vs \ [] \ es) \mid \mathbf{new} \ C(vs \ [] \ es)$
v	$::= \mathbf{new} \ C(vs)$
Γ	$::= x_1 : C_1 \dots x_n : C_n$
S	$::= \dots \textit{e.g.} \ \mathbf{Pre} : P \ \mathbf{Post} : P$
P	$::= \dots \textit{e.g.} \ \text{First order logic}$

Fig. 1. Syntax of the trait system

We check if methods are well-typed with judgments of form $Ds; Name \vdash M : OK$. This judgment can be read as: in the definition table, the method M defined under the definition $Name$ is correct. The typing rules of Fig. 2 are explained in the technical report [41] in detail. The first four rules type different expressions and collect the information of these expressions to prove with rule MOK that a method fulfills its specification. In the rule MOK with keyword **verify**, we call a verifier to prove each method once. Abstract methods (ABSOK) are always correct. Rule BODYOK ensures that all methods in a body are correctly typed.

3.3 Flattening Semantics

When we implement methods in several traits, we have to check that these traits are compatible when they are composed. This process to derive a complete class from a set of traits is called flattening. We follow the traditional flattening semantics [20]. A class that is defined by composing several traits is obtained by flattening rules. All methods are direct members of the class [20]. Overall, our flattening process works as a big step reduction arrow, where we reduce a trait expression into a well-typed and verified body.

To introduce our flattening rules in Fig 3, we first define the helper functions. The function *allMeth* collects all method headers with the same name as m in all input bodies (Definition 1). When two *Bodys* are composed (Definition 2), the implemented interfaces are united and the methods are composed. The composition of methods (Definition 3) collects methods that are only defined in one of the input sets. If a method is in both sets, it is composed (Definition 4). Here, we distinguish four cases. If one method is abstract and the other is concrete, we have to show that the precondition of the abstract method implies the precondition of the concrete method. Additionally, the postcondition of the concrete one has to imply the postcondition of the abstract one. This is similar to Liskov's

$$\begin{array}{c}
 \frac{}{\Gamma \vdash x : \Gamma(x) \dashv \text{result} : \Gamma(x) \ \& \ \text{result} = x \models \text{true}} \quad (\text{X}) \\
 \\
 \frac{
 \begin{array}{l}
 S \text{ method } C \ m(C_1 \ x_1 \dots C_n \ x_n); \in \ \text{methods}(C_0) \\
 \Gamma \vdash e_0 : C_0 \dashv P_0 \models P'_0 \dots \Gamma \vdash e_n : C_n \dashv P_n \models P'_n \\
 x'_0 \dots x'_n \ \text{fresh} \quad S' = S[\text{this} := x'_0, x_1 := x'_1, \dots, x_n := x'_n] \\
 P = (\text{result} : C \ \& \ P_0[\text{result} := x'_0] \ \& \ \dots \ \& \ P_n[\text{result} := x'_n] \\
 \quad \& \ (Pre(S') \implies Post(S')))
 \end{array}
 }{\Gamma \vdash e_0.m(e_1 \dots e_n) : C \dashv P \models P'_0 \ \& \ \dots \ \& \ P'_n \ \& \ Pre(S')} \quad (\text{METHOD}) \\
 \\
 \frac{
 \begin{array}{l}
 \Gamma \vdash e_1 : C_1 \dashv P_1 \models P'_1 \dots \Gamma \vdash e_n : C_n \dashv P_n \models P'_n \\
 \text{getters}(C) = S_1 \text{ method } C_1 \ x_1(); \dots S_n \text{ method } C_n \ x_n(); \quad x'_1 \dots x'_n \ \text{fresh} \\
 S'_i = S_i[\text{this} := \text{result}] \quad P'_i = (P_i[\text{result} := x'_i] \ \& \ (Pre(S'_i) \implies \text{result}.x_i() = x'_i)) \\
 P = (\text{result} : C \ \& \ P'_1 \ \& \ \dots \ \& \ P'_n)
 \end{array}
 }{\Gamma \vdash \text{new } C(e_1 \dots e_n) : C \dashv P \models P'_1 \ \& \ \dots \ \& \ P'_n \ \& \ Pre(S'_1) \ \& \ \dots \ \& \ Pre(S'_n)} \quad (\text{NEW}) \\
 \\
 \frac{
 \Gamma \vdash e : C' \dashv P \models P' \quad C' \ \text{instance of } C
 }{\Gamma \vdash e : C \dashv P \models P'} \quad (\text{SUB}) \\
 \\
 \frac{
 \Gamma = \text{this} : \text{Name}, \ x_1 : C_1, \ \dots, \ x_n : C_n \quad \Gamma \vdash e : C \dashv P \models P' \\
 \text{verify } Ds \vdash (\Gamma \ \& \ Pre(S) \ \& \ P) \models (P' \ \& \ Post(S))
 }{Ds; \ \text{Name} \vdash S \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n) \ e; : \ \text{OK}} \quad (\text{MOK}) \\
 \\
 \frac{}{Ds; \ \text{Name} \vdash S \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n); : \ \text{OK}} \quad (\text{AbsMOK}) \\
 \\
 \frac{
 \text{Body} = \{\text{interface? } [Cs] \ M_1 \dots M_n\} \\
 Ds; \ \text{Name} \vdash M_1 : \ \text{OK} \ \dots \ Ds; \ \text{Name} \vdash M_n : \ \text{OK}
 }{Ds; \ \text{Name} \vdash \text{Body} : \ \text{OK}} \quad (\text{BODYTYPED})
 \end{array}$$

Fig. 2. Expression typing rules

substitution principle [34]. The second case is the symmetric variant of the first case. In the third and fourth case, two abstract methods are composed. Here, the specification of one abstract method has to imply the specification of the other abstract method such that an implementation can still satisfy all specifications of abstract methods. If both method are concrete, the composition is correctly left undefined. This composition error can be resolved by making one method m abstract in the *Body*, as defined in Definition 5. The resulting *Body* is similar with the difference that the implementation of the method m is omitted. The flattening rules in Fig. 3 are explained in the following in detail. In these rules, a set of traits is flattened to a declaration containing all methods. If abstract and concrete methods with the same name are composed, Definitions 2-4 are used to guarantee correctness of the composition.

Definition 1 (All Methods). $allMeth(m, \text{Body}) = \{MH; \mid \text{Body} \in \text{Body}, \text{Body}(m) = MH; \}$

Definition 2 (Body Composition). $\text{Body}_1 + \text{Body}_2 = \text{Body}$
 $\{\text{interface? } [Cs_1] \ Ms_1\} + \{\text{interface? } [Cs_2] \ Ms_2\} =$
 $\{\text{interface? } [Cs_1 \cup Cs_2] \ Ms_1 + Ms_2\}$

$$\begin{array}{c}
\frac{D'_1 \dots D'_n \vdash D_1 \Downarrow D'_1 \quad \dots \quad D'_1 \dots D'_n \vdash D_n \Downarrow D'_n}{D_1 \dots D_n \Downarrow D'_1 \dots D'_n} \text{ (FLATTOP)} \\
\\
\frac{Ds; \text{Name} \vdash E \Downarrow \text{Body} \quad \text{if Name of form } C \text{ then } \text{abs}(\text{Body}) = S \ T \ x_1(); \dots \ S \ T \ x_n();}{Ds \vdash \text{Name} = E \Downarrow \text{Name} = \text{Body}} \text{ (DFLAT)} \\
\\
\frac{\begin{array}{l} \text{Body} = \{\text{interface? } [Cs] \ M_1 \dots M_n \} \\ \text{Body}' = \{\text{interface? } [Cs] \ M_1 \dots M_n \ Ms \} \\ Ms = \{\Sigma \text{allMeth}(Ds, Cs, m) \mid m \in \text{dom}(Cs) \text{ and } m \notin \text{dom}(\text{Body})\} \\ Ds; \text{Name} \vdash \text{Body}' : \text{OK} \end{array}}{Ds; \text{Name} \vdash \text{Body} \Downarrow \text{Body}'} \text{ (BFLAT)} \\
\\
\frac{}{Ds; \text{Name} \vdash t \Downarrow Ds(t)} \text{ (TFLAT)} \\
\\
\frac{Ds; \text{Name} \vdash E_1 \Downarrow \text{Body}_1 \quad Ds; \text{Name} \vdash E_2 \Downarrow \text{Body}_2}{Ds; \text{Name} \vdash E_1 + E_2 \Downarrow \text{Body}_1 + \text{Body}_2} \text{ (+FLAT)} \\
\\
\frac{Ds; \text{Name} \vdash E \Downarrow \text{Body} \quad \text{Body} = \{[Cs] \ \overline{M}_1 \ S \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n); \ \overline{M}_2\} \\ \text{Body}' = \{[Cs] \ \overline{M}_1 \ S \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n); \ \overline{M}_2\}}{Ds; \text{Name} \vdash E[\text{makeAbstract } m] \Downarrow \text{Body}'} \text{ (ABSFLAT)}
\end{array}$$

Fig. 3. Flattening rules

Definition 3 (Methods Composition). $Ms_1 + Ms_2 = Ms$

- $(M \ Ms_1) + Ms_2 = M (Ms_1 + Ms_2)$
if $\text{methName}(M) \notin \text{dom}(Ms_2)$
- $(M_1 \ Ms_1) + (M_2 \ Ms_2) = M_1 + M_2 (Ms_1 + Ms_2)$
if $\text{methName}(M_1) = \text{methName}(M_2)$
- $\emptyset + Ms = Ms$

Definition 4 (Method Composition). $M_1 + M_2 = M$

- $S \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n) \ e; + S' \ \text{method } C \ m(C_1 \ \dots C_n \ -);$
= $S \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n) \ e;$
if $\text{Pre}(S')$ implies $\text{Pre}(S)$ and $\text{Post}(S)$ implies $\text{Post}(S')$
- $MH_1; + MH_2 \ e; = MH_2 \ e; + MH_1;$
- $S \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n); + S' \ \text{method } C \ m(C_1 \ \dots C_n \ -);$
= $S \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n);$
if $\text{Pre}(S')$ implies $\text{Pre}(S)$ and $\text{Post}(S)$ implies $\text{Post}(S')$
- $S \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n); + S' \ \text{method } C \ m(C_1 \ \dots C_n \ -);$
= $S' \ \text{method } C \ m(C_1 \ x_1 \dots C_n \ x_n);$
if $(\text{Pre}(S) \text{ implies } \text{Pre}(S') \text{ and } \text{Post}(S') \text{ implies } \text{Post}(S))$
and not $(\text{Pre}(S') \text{ implies } \text{Pre}(S) \text{ and } \text{Post}(S) \text{ implies } \text{Post}(S'))$

Definition 5 (Body Abstraction). $\text{Body}[\text{makeAbstract } m]$

$$\begin{aligned}
& \{[Cs] \ Ms_1 \ S \ \text{method } C \ m(Cxs); \ Ms_2\}[\text{makeAbstract } m] \\
& = \{[Cs] \ Ms_1 \ S \ \text{method } C \ m(Cxs); \ Ms_2\}
\end{aligned}$$

FLATTOP. The first rule flattens a set of declarations $D_1 \dots D_n$ to a set $D'_1 \dots D'_n$. We express this rule in a non-computational way: we assume to know

the resulting $D'_1 \dots D'_n$, and we use them as a guide to compute them. Note that if there is a resulting $D'_1 \dots D'_n$ then it is unique; flattening is a deterministic process and $D'_1 \dots D'_n$ are used only to type check the results. They are not used to compute the shape of the flattened code.

Non computational rules like this are common with nominal type systems [27] where the type signatures of all classes and methods can be extracted before the method bodies are verified.

DFLAT. This rule flattens an individual definition by flattening the trait expression. When the flattening produces a class definition, we also check that the body denotes an instantiable class; a class whose only abstract methods are valid getters. The function $abs(Body)$ returns the abstract methods.

BFLAT It may look surprising that the $Body$ does not flatten to itself. This represents what happens in most programming languages, where implementing an interface implicitly imports the abstract signature for all the methods of that interface. In the context of verification also the specification of such interface methods is imported. In concrete, $Body'$ is like $Body$, but we add Ms by collecting all the methods of the interfaces that are not already present in the $Body$.

Moreover, we check that all the methods defined in the class respect the typing and the specification defined in the interfaces: if a class has S **method** `Foo foo()`; or S **method** `Foo foo() e`; and there is a S' **method** `Foo foo()`; in the interface, then S must respect the specification S' . The system then checks that the $Body$ is well-typed and verified by calling Ds ; $Name \vdash M_i : OK$

TFLAT. A trait t is flattened to its declaration $Ds(t)$.

+FLAT. The composition of two expression E_1 and E_2 , where both expressions are first reduced to $Body_1$ and $Body_2$, results in the composition of these bodies as defined in Definition 2.

ABSFLAT. An expression E where one method m is made abstract flattens to a $Body'$. We know that E flattens to $Body$. The only difference between $Body$ and $Body'$ is that the one method m is abstract in $Body'$. In $Body$, the method can be abstract or concrete.

3.4 Soundness of the Trait-based CbC Process

In this section, we formulate our main result of the TraitCbC process. We prove soundness of the flattening process with a parametric logic. The proofs of the lemmas and theorems are in the technical report [41]. We claim that if you have a language without code reuse and with sound and modular PhV verification then the language supports CbC simply by adding traits to the language. That is, traits intrinsically enable a CbC program construction process.

To prove soundness of the refinement process of TraitCbC (Theorem 2: Sound CbC Process) as exemplified in Section 2, we have to show that the flattening process is correct (Theorem 1: General Soundness). In turn, to prove General Soundness, we need two lemmas which state that the composition of traits is correct (Lemma 1) and that a trait after the `makeAbstract` operation is still correct (Lemma 2).

In Lemma 1, we have well-typed definitions Ds , and two well-typed and verified traits in Ds , and the resulting trait/class is also well-typed and verified.

Lemma 1 (Composition correct).

If $Ds(t_1) = Body_1$, $Ds(t_2) = Body_2$, $Ds(Name) = Body$, $Ds; t_1 \vdash Body_1 : OK$,
 $Ds; t_2 \vdash Body_2 : OK$, and $Body_1 + Body_2 = Body$,
 then $Ds; Name \vdash Body : OK$

Lemma 2 shows that if we have a well-typed and verified trait, the operation `makeAbstract` results in a trait/class that is also well-typed and verified.

Lemma 2 (MakeAbstract correct).

If $Ds(t) = Body$, $Ds(Name) = Body'$, $Ds; t \vdash Body : OK$,
 and $Body[\text{makeAbstract } m] = Body'$,
 then $Ds; Name \vdash Body' : OK$

With these Lemmas, we can prove Theorem 1. Given a sound and modular verification language, then all programs that flatten are well-typed and verified. In a modular verification language, a method can be fully verified using only the information contained in the method declaration and the specification of any used method. Moreover, our parametric logic must support at least a commutative and associative *and* (but of course other ways to merge knowledge could work too) and a transitive *implication* (but of course other forms of logical consequence could work too).

Theorem 1 (General Soundness).

For all programs Ds where Ds flattens to Ds' , and Ds' is well-typed;
 that is, for all $Name = Body \in Ds'$, we have $Ds'; Name \vdash Body : OK$.

We now show that the `TraitCbC` process is sound. Theorem 2 states that starting with one abstract method and a set of verified traits, the composed program is also verified.

Theorem 2 (Sound CbC Process).

Starting from a fully abstract specification t_0 , and some refinement steps $t_1 \dots t_n$, we can write $C = t_0 + \dots + t_n$ as our whole `CbC` refinement process; where $t_0 + t_1$ is the application of the first refinement step. If we use `CbC` to construct programs, we can start from verified atomic units and get a verified result. Formally, if

$$\begin{array}{l} t_0 = \{MH\} \quad t_1 = \{Ms_1\} \quad \dots \quad t_n = \{Ms_n\} \quad \text{are well-typed, and} \\ t_0 = \{MH\} \qquad \qquad \qquad t_0 = \{MH\} \\ t_1 = \{Ms_1\} \quad \dots \quad t_n = \{Ms_n\} \quad \Downarrow \quad t_1 = \{Ms_1\} \quad \dots \quad t_n = \{Ms_n\} \\ C = t_0 + \dots + t_n \qquad \qquad \qquad C = Body \end{array}$$

then $C = Body$ is well-typed.

Proof. This is a special case of Theorem 1.

Theorem 2 shows clearly that trait composition intrinsically enables a `CbC` refinement process: A object-oriented programming language with traits and a corresponding specification language supports an incremental `CbC` approach.

	Classic CbC	TraitCbC
Language	Additional rules for a programming language.	Programming language with traits. Needs specification language.
Tool support	Pen and paper. Some specialized tools available.	Relies on prevalent PhV verification tools.
Construction Rules	Specific refinement rules.	Refinement by composition of traits.
Debugging	Guarantees the correctness of each refinement step. Only refinements without abstract statement are directly verified.	Guarantees the correctness of each refinement step. Each method is specified such that each refinement can directly be verified.
Proof complexity	Many, but small proofs.	Any granularity of proofs.
Reuse	Refinement steps cannot be reused; only fully implemented methods can.	Each verified method in a trait can be reused.
Applications	Focuses on small but correctness-critical algorithms.	As TraitCbC is based on PhV, it can be used in areas of PhV. Additionally, traits are beneficial for incremental development approaches and development of software product lines.

Table 1. Comparison of TraitCbC with classical CbC

4 Trait-based Correctness-by-Construction in Comparison to Classical CbC

In this section, we discuss the benefits of TraitCbC in comparison to classical CbC. To do this, we describe classical CbC first.

Classical correctness-by-construction (CbC) [19, 30, 37] is an incremental approach to construct programs. CbC uses a *Hoare triple* specification $\{P\} S \{Q\}$ stating that if the precondition P holds, and the statement S is executed, then the statement terminates and postcondition Q holds. The CbC refinement process starts with a Hoare triple where the statement S is abstract. This abstract statement can be seen as a hole in the program that needs to be filled. With a set of refinement rules, an abstract statement is replaced by more concrete statements (i.e., statements in the guarded command language [19] that can contain further abstract statements). The process stops, when all abstract statements are refined to concrete statements so that no holes remain in the program. As each refinement rule is sound and each correct application of a refinement rule guarantees to satisfy the starting Hoare triple, the resulting program is correctness-by-construction [30]. The CbC process is strictly tied to a set of predefined refinement rules. A programmer cannot deviate from this concept. To apply a refinement rule, it has to be checked that conditions of the rule application are satisfied. This is done by pen-and-paper or with specialized tools [42].

In Table 1, we compare TraitCbC and classical CbC:

Language The classical CbC approach is external to a programming language. It needs the definition of refinement rules. TraitCbC is usable with languages that have traits, a specification language, and a corresponding verification framework. In this work, we focus on object-orientation, but the general TraitCbC programming guideline presented in this paper is also suitable for functional programming

environments using abstract and concrete functions with specifications instead of traits and methods.

Tool Support To use one of the approaches, tool support is desired. For classical CbC, mostly pen and paper is used. There are a few specialized tools such as CorC [42], tool support for ArcAngel [38], and SOCOS [4, 5]. These tools force a certain programming procedure on the user. This procedure can be in conflict with their preferred programming style. For TraitCbC, tools for post-hoc verification can be reused. There are tools for many languages such as Java [3], C [16], C# [7, 8]. Other languages are integrated with their verifier from the start, e.g., Spec# [8] and Dafny [32]. TraitCbC as presented in this paper is a core calculus, designed to show the feasibility of the concept. We believe that scaling up TraitCbC to a complete programming language reusing existing verification techniques would be feasible and would result in a similarly expressive verification process, but supporting more flexible program composition. In Section 5, we show how a prototype can be constructed by using the KeY verifier [3].

Construction Rules To construct a program, classical CbC has a strict concept of refinement rules. A programmer cannot deviate from the granularity of the rules. In contrast, PhV does not give a mandatory guideline how to construct programs. TraitCbC is a bridge between both extremes. Programs can be constructed stepwise as with classical CbC, but if desired, any number of refinement steps can be condensed up to PhV based programming.

Debugging If errors occur in the development process, TraitCbC gives early and detailed information. By specifying the method under development and any abstract method that is called by this method, we can directly verify the correctness of the method under development. We assume that the introduced abstract methods will be correctly implemented in further refinement steps. With each step, the programmer gets closer to the solution until finally all abstract methods are implemented. Classical CbC relies on the same process, but here the abstract statements (similar to our abstract methods) are not explicitly specified by the programmer. Additional specifications in classical CbC are introduced only with some rules such as an intermediate condition in the composition rule. Then, these specifications are propagated through the program to be constructed. When arriving at a leaf in the refinement process, the correctness of the statement can be guaranteed. The problem in classical CbC is that all refinement steps where abstract statements occur cannot be verified directly. In the worst case, a wrong specification is found only after a few refinement steps.

Proof Complexity TraitCbC can have the same granularity and also the same proof effort as classical CbC, since each method implementation can correspond to just one refinement step. The advantage of TraitCbC is that programmers can freely implement a method body. They must not stick to the same granularity as in the classical CbC refinement rules. As in PhV, they can implement a complete method in one step. The programmer can balance proof complexity against verifier calls.

Reuse If we want to reuse developed methods or refinement steps, the approaches differ. In classical CbC, no refinement steps can be reused. A fully refined method

can be reused in both approaches. For TraitCbC, we can easily reuse even very small units of code, since they are represented as methods in the traits.

Applications The classical CbC approach does not scale well to development procedures for complete software system. Rather, individual algorithms can be developed with CbC [50]. As soon as we scale TraitCbC to real languages, we have the same application scenarios as PhV. As argued by Damiani et al. [18] traits enable an incremental process of specifying and verifying software. Bettini et al. [10] proposed to use traits for software product line development and highlighted the benefits of fine-grained reuse mechanisms. Here, TraitCbC’s guideline is suitable for constructing new product lines step by step from the beginning.

Summary In summary, TraitCbC bridges the gap between PhV and CbC. It enables a CbC process for trait-based languages without introducing refinement rules. The concrete realization of specifying and verifying methods is similar to PhV, but additionally to PhV, TraitCbC provides an incremental development process. This development process combined with the flexibility of traits allows correct methods to be developed in small and reusable steps. Moreover, we have introduced a core calculus and proved that the construction and composition of trait-based programs is correct.

5 Proof-of-Concept Implementation

In this section, we describe the implementation, which instantiates TraitCbC in Java with JML [31] as specification language and KeY [3] as verifier for Java code. Our trait implementation is based on interfaces with default implementation. Our open source tool is implemented in Java and integrated as plug-in in the Eclipse IDE.⁸ Besides this prototype, other languages with a suitable verifier, such as Dafny [32] and OpenJML [17], can also be used to implement TraitCbC.

In Listing 1, we show the concrete syntax. Each method in a trait is specified with JML with the keywords `requires` and `ensures` for the pre- and postcondition. To verify the correctness of programs, we need two steps. First, we verify the correctness of a method implemented in a trait w.r.t. its specification. Second, for trait composition, our implementation checks the correct composition for all methods (cf. Definition 2). It is verified that the specification of a concrete method satisfies the specification of the abstract one with the same signature (cf. Definition 4). These verification goals are sent to KeY, which starts an automatic verification attempt. The syntax of trait composition is shown in line 24. In a separate tc-file, the name of the resulting trait is given and the composed traits are connected with a plus operator.

```

1 public interface MaxElement1 {
2   /*@ requires list.size() > 0;
3     @ ensures (\forall int n; list.contains(n);
4     @ \result >= n) & list.contains(\result);
5     @*/

```

⁸ Tool and evaluation at <https://github.com/TUBS-ISF/CorC/tree/TraitCbC>

```

6   public default int maxElement(List list) {
7       if (list.size() == 1) return accessHead(list);
8       if (list.element() >= maxElement(list.tail()))
9           { return accessHead(list) }
10          else { return maxTail(list) } }
11
12  /*@ requires list.size() > 0;
13     @ ensures \result == list.element();
14     @*/
15  public int accessHead(List list);
16
17  /*@ requires list.size() > 1;
18     @ ensures (\forall int n; list.tail().contains(n);
19     @ \result >= n) & list.tail().contains(\result);
20     @*/
21  public int maxTail(List list);
22  }
23
24  ComposedMax = MaxElement1 + MaxElement2

```

Listing 1. Example in our implementation

Evaluation We evaluate our implementation by a feasibility study. First, we reimplemented an already verified case study in our trait-based language. We used the IntList [43] case study, which is a small software product line (SPL) with a common code base and several features extending this code base. Here, we can show that our trait-based language also facilitates reuse. The IntList case study implements functionality to insert integers to a list in the base version. Extensions are the sorting of the list and different insert options (e.g., front/back). We implement five methods that exists in different variants with our trait-based CbC approach. We implement the case study in different granularities. The coarse-grained version is similar to the SPL implementation we started with [43], confirming that traits are also amenable to implement SPLs as shown by Bettini et al. [10]. The fine-grained version implements the five methods incrementally with 12 refinement steps. We can reuse 6 of these steps during the construction of method variants.

We also implement three more case studies BankAccount [48], Email [25], and Elevator [39] with TraitCbC and CbC to show that it is feasible to implement object-oriented programs with both approaches. We used CorC [42] as an instance of a CbC tool. We were able to implement nine classes and verify 34 methods with a size of 1-20 lines of code. For future work, a user study is necessary to evaluate the usability of TraitCbC in comparison to CbC to confirm our stated advantages.

6 Related Work

Traits are introduced in many languages to support clean design and reuse, for example Smalltalk [20], Java [12] by utilizing default methods in interfaces,

and other Java-like languages [11, 33, 45]. The trait language TraitRecordJ was extended to support post-hoc verification of traits [18]. The authors added specifications of methods in traits for the verification of correct trait composition and proposed a modular and incremental verification process. None of these trait languages were used to formulate a refinement process to create correct programs. They only focus on code reuse or post-hoc verification.

Automatic verification is widely used for different programming languages. The object-oriented language Eiffel focuses on design-by-contract [35, 36]. All methods in classes are specified with pre-/postconditions and invariants for verification purposes. The tool AutoProof [29, 49] is used to verify the correctness of implemented methods. It translates methods to logic formulas, and an SMT solver proves the correctness. For C#, programs written in the similar language Spec# [8] are verified with Boogie. That is, code and specification are translated to an intermediate language and verified [7]. For C, the tool VCC [16] reuses the Spec# tool chain to verify programs. The tool VeriFast [28] is able to verify C and Java programs specified in separation logic. For Java, KeY [3] and OpenJML [17] verify programs specified with JML. TraitCbC is parametric in the specification language, meaning that a trait-based language with a specification language and a corresponding program verifier can be used to instantiate TraitCbC. In our implementation, we use KeY [3] to prove the correctness of methods and trait composition.

Event-B [1] is a related correctness-by-construction approach. In Event-B, automata-based systems are specified and refined to a concrete implementation. Event-B is implemented in the Rodin platform [2]. In comparison to CbC by Kourie and Watson [30] as used in this paper, Event-B works on a different abstraction level with automata-based systems instead of program code. The CbC approaches of Back et al. [6] and Morgan [37] are also related. Back et al. [6] start with explicit invariants and pre-/postconditions to refine an abstract program to a concrete implementation, while Kourie and Watson only start with a pre-/postcondition specification. These refinement approaches use specific refinement rules to construct programs which are external to the programming language. With TraitCbC, we propose a refinement procedure that is part of the language by using trait composition.

Abstract execution [47] verifies the correctness of methods with abstract, but formally specified expressions. Abstract Execution is similar to our refinement procedure where abstract methods are called in methods under construction. The difference is that abstract execution extends a programming language to use any expression in the abstract part, not only method calls. Therefore, abstract execution can better reason about irregular termination (e.g., break/continue) of methods. In comparison to TraitCbC, abstract execution is a verification-centric approach without a guideline on how to construct programs.

Synthesis of function summaries is also related [26, 14, 44]. Here, verification tools automatically synthesize pre-/postconditions from functions to achieve modular verification and speed up the verification time. In comparison, TraitCbC is a complete software development approach where specification and code are

developed simultaneously by a developer to achieve a correct solution. Function summaries are just a verification technique.

7 Conclusion

In this work, we present TraitCbC that guides programmers to correct implementations. In comparison to classical CbC, TraitCbC uses method calls and trait composition instead of refinement rules to guarantee functional correctness. We formalize the concept of a trait-based object-oriented language where the specification language is parametric to allow a broader range of languages to adopt this concept. The main advantage of TraitCbC is the simplicity of the refinement process that supports code and proof reuse.

As future work, we want to investigate how TraitCbC can be used to construct software product lines. As proposed by Bettini et al. [10], trait languages are able to implement SPLs. We want to extend the guideline of TraitCbC to construct SPLs with a refinement-based procedure that guarantees the correctness of the whole SPL. To reduce specification effort in TraitCbC, inheritance of traits is useful. Another option is to integrate the concept of Rebêlo et al. [40] which supports the design-by-contract approach with AspectJML and integrates crosscutting contract modularization to reduce redundant specifications.

Since TraitCbC is parametric in the specification logic, TraitCbC’s soundness only holds if such logic is consistent when composed in the presented manner. In particular, the logic needs to take into account ill-founded specifications and non-terminating recursion. In verification, ill-founded specifications and termination issues are often considered as a second step⁹, separately from the verification of individual methods, and our prototype still does not yet take care of this second step. That means that methods are verified under the assumption that all other methods respect their contracts. If ill-founded specifications and non-terminating recursion are handled naively, verification might be unsound because of ill-founded reasoning. The technical report [41] shows that this problem is even more pervasive in the case of trait composition or any other form of multiple inheritance: naive composition of correct traits may produce incorrect results.

References

1. Abrial, J.: Modeling in Event-B - System and Software Engineering. Cambridge University Press (2010)
2. Abrial, J.R., Butler, M., Hallerstede, S., Hoang, T.S., Mehta, F., Voisin, L.: Rodin: An Open Toolset for Modelling and Reasoning in Event-B. STTT 12(6), 447–466 (2010)
3. Ahrendt, W., Beckert, B., Bubel, R., Hähnle, R., Schmitt, P.H., Ulbrich, M.: Deductive Software Verification—The KeY Book: From Theory to Practice, vol. 10001. Springer (2016)

⁹ For example, Dafny approximately checks that the functions used in a specification form an acyclic graph.

4. Back, R.J.: Invariant Based Programming: Basic Approach and Teaching Experiences. *FAOC* 21(3), 227–244 (2009)
5. Back, R.J., Eriksson, J., Myreen, M.: Testing and Verifying Invariant Based Programs in the SOCOS Environment. In: *TAP. LNCS*, vol. 4454, pp. 61–78. Springer (2007)
6. Back, R.J., Wright, J.: *Refinement Calculus: A Systematic Introduction*. Springer Science & Business Media (2012)
7. Barnett, M., Fähndrich, M., Leino, K.R.M., Müller, P., Schulte, W., Venter, H.: Specification and Verification: The Spec# Experience. *Communication of the ACM* 54(6), 81–91 (Jun 2011)
8. Barnett, M., Leino, K.R.M., Schulte, W.: The Spec# Programming System: An Overview. In: *CASSIS*. pp. 49–69. Springer (2004)
9. ter Beek, M.H., Cleophas, L., Schaefer, I., Watson, B.W.: X-by-Construction. In: *ISoLA*. pp. 359–364. Springer International Publishing, Cham (2018)
10. Bettini, L., Damiani, F., Schaefer, I.: Implementing Software Product Lines Using Traits. In: *SAC*. pp. 2096–2102 (2010)
11. Bettini, L., Damiani, F., Schaefer, I., Stocco, F.: TRAITRECORDJ: A programming Language with Traits and Records. *Science of Computer Programming* 78(5), 521–541 (2013)
12. Bono, V., Mensa, E., Naddeo, M.: Trait-Oriented Programming in Java 8. In: *PPPJ*. pp. 181–186 (2014)
13. Chapman, R.: Correctness by Construction: A Manifesto for High Integrity Software. In: *SCS*. pp. 43–46 (2006)
14. Chen, H.Y., David, C., Kroening, D., Schrammel, P., Wachter, B.: Synthesising interprocedural bit-precise termination proofs (t). In: *ASE*. pp. 53–64 (2015)
15. Clements, P., Northrop, L.: *Software Product Lines: Practices and Patterns*. Addison-Wesley (2002)
16. Cohen, E., Dahlweid, M., Hillebrand, M., Leinenbach, D., Moskal, M., Santen, T., Schulte, W., Tobies, S.: VCC: A Practical System for Verifying Concurrent C. In: *TPHOLS. LNCS*, vol. 5674, pp. 23–42. Springer (2009)
17. Cok, D.R.: OpenJML: JML for Java 7 by Extending OpenJDK. In: *NFM. LNCS*, vol. 6617, pp. 472–479. Springer (2011)
18. Damiani, F., Dovland, J., Johnsen, E.B., Schaefer, I.: Verifying Traits: An Incremental Proof System for Fine-Grained Reuse. *FAOC* 26(4), 761–793 (2014)
19. Dijkstra, E.W.: *A Discipline of Programming*. Prentice Hall (1976)
20. Ducasse, S., Nierstrasz, O., Schärli, N., Wuyts, R., Black, A.P.: Traits: A Mechanism for Fine-Grained Reuse. *TOPLAS* 28(2), 331–388 (2006)
21. Flatt, M., Krishnamurthi, S., Felleisen, M.: Classes and Mixins. In: *POPL*. pp. 171–183 (1998)
22. Gries, D.: *The Science of Programming*. Springer (1987)
23. Haftmann, F., Krauss, A., Kunčar, O., Nipkow, T.: Data Refinement in Isabelle/HOL. In: *ITP*. pp. 100–115. Springer (2013)
24. Hall, A., Chapman, R.: Correctness by Construction: Developing a Commercial Secure System. *Software, IEEE* 19(1), 18–25 (Jan 2002)
25. Hall, R.J.: Fundamental Nonmodularity in Electronic Mail. *ASE* 12(1), 41–79 (2005)
26. Hoare, C.A.R.: Procedures and parameters: An axiomatic approach. In: *Symposium on Semantics of Algorithmic Languages*. pp. 102–116. Springer (1971)
27. Igarashi, A., Pierce, B.C., Wadler, P.: Featherweight Java: A Minimal Core Calculus for Java and GJ. *TOPLAS* 23(3), 396–450 (2001)

28. Jacobs, B., Smans, J., Piessens, F.: A Quick Tour of the VeriFast Program Verifier. In: APLAS. LNCS, vol. 6461, pp. 304–311. Springer (2010)
29. Khazeev, M., Rivera, V., Mazzara, M., Johard, L.: Initial Steps Towards Assessing the Usability of a Verification Tool. In: SEDA. Advances in Intelligent Systems and Computing, vol. 717, pp. 31–40. Springer (2016)
30. Kourie, D.G., Watson, B.W.: The Correctness-by-Construction Approach to Programming. Springer Science & Business Media (2012)
31. Leavens, G.T., Baker, A.L., Ruby, C.: JML: a Java Modeling Language. In: Formal Underpinnings of Java Workshop (at OOPSLA’98). pp. 404–420. Citeseer (1998)
32. Leino, K.R.M.: Dafny: An Automatic Program Verifier for Functional Correctness. In: LPAR. pp. 348–370. Springer (2010)
33. Liquori, L., Spiwack, A.: FeatherTrait: A Modest Extension of Featherweight Java. TOPLAS 30(2), 1–32 (2008)
34. Liskov, B.H., Wing, J.M.: A Behavioral Notion of Subtyping. TOPLAS 16(6), 1811–1841 (1994)
35. Meyer, B.: Eiffel: A Language and Environment for Software Engineering. JSS 8(3), 199–246 (1988)
36. Meyer, B.: Applying “Design by Contract”. Computer 25(10), 40–51 (1992)
37. Morgan, C.: Programming from Specifications. Prentice Hall, 2nd edn. (1994)
38. Oliveira, M.V.M., Cavalcanti, A., Woodcock, J.: ArcAngel: A Tactic Language for Refinement. FAOC 15(1), 28–47 (2003)
39. Plath, M., Ryan, M.: Feature Integration Using a Feature Construct. Science of Computer Programming 41(1), 53–84 (2001)
40. Rebêlo, H., Leavens, G.T., Bagherzadeh, M., Rajan, H., Lima, R., Zimmerman, D.M., Cornélio, M., Thüm, T.: AspectJML: Modular Specification and Runtime Checking for Crosscutting Contracts. In: MODULARITY. p. 157–168. ACM, New York, NY, USA (2014)
41. Runge, T., Potanin, A., Thüm, T., Schaefer, I.: Traits for correct-by-construction programming (2022)
42. Runge, T., Schaefer, I., Cleophas, L., Thüm, T., Kourie, D., Watson, B.W.: Tool Support for Correctness-by-Construction. In: FASE. pp. 25–42. Springer (2019)
43. Scholz, W., Thüm, T., Apel, S., Lengauer, C.: Automatic Detection of Feature Interactions Using the Java Modeling Language: An Experience Report. In: SPLC. ACM, New York, NY, USA (2011)
44. Sery, O., Fedyukovich, G., Sharygina, N.: Interpolation-Based Function Summaries in Bounded Model Checking. In: HVC. pp. 160–175. Springer (2012)
45. Smith, C., Drossopoulou, S.: Chai: Traits for Java-Like Languages. In: ECOOP. pp. 453–478. Springer, Berlin, Heidelberg (2005)
46. Sozeau, M., Oury, N.: First-Class Type Classes. In: TPHOLs. pp. 278–293. Springer (2008)
47. Steinhöfel, D., Hähnle, R.: Abstract Execution. In: FM. pp. 319–336. Springer (2019)
48. Thüm, T., Schaefer, I., Apel, S., Hentschel, M.: Family-Based Deductive Verification of Software Product Lines. In: GPCE. p. 11–20. ACM (2012)
49. Tschannen, J., Furia, C.A., Nordio, M., Polikarpova, N.: AutoProof: Auto-Active Functional Verification of Object-Oriented Programs. In: TACAS. LNCS, vol. 9035, pp. 566–580. Springer (2015)
50. Watson, B.W., Kourie, D.G., Schaefer, I., Cleophas, L.: Correctness-by-Construction and Post-hoc Verification: A Marriage of Convenience? In: ISoLA. LNCS, vol. 9952, pp. 730–748. Springer (2016)