

Instance Pointcuts

Selecting Object Sets Based on Their Usage History

Christoph Bockisch

TRESE, University of Twente, 7500AE
Enschede, The Netherlands
c.m.bockisch@utwente.nl

Kardelen Hatun

TRESE, University of Twente, 7500AE
Enschede, The Netherlands
k.hatun@utwente.nl

Mehmet Aksit

TRESE, University of Twente, 7500AE
Enschede, The Netherlands
m.aksit@utwente.nl

Abstract

At runtime, how objects have to be handled frequently depends on how they were used before. But with current programming-language support, selecting objects according to their previous usage patterns often results in scattered and tangled code. In this study, we propose a new kind of pointcut, called *Instance Pointcuts*, for maintaining sets that contain objects with a specified usage history. Instance pointcut specifications can be reused, by refining their selection criteria, e.g., by restricting the scope of an existing instance pointcut; and they can be composed, e.g., by set operations. These features make instance pointcuts easy to evolve according to new requirements. Our approach improves modularity by providing a fine-grained mechanism and a declarative syntax to create and maintain usage-specific object sets.

Categories and Subject Descriptors D.3.3 [Language Constructs and Features]: [classes and objects]

General Terms Languages, Design

Keywords programming languages, aspect-oriented programming, instance pointcuts, objects

1. Introduction

In object-oriented programming (OOP), the encapsulated state and the provided behavior of objects is dictated by their type. Nevertheless, often objects of the same type need to be treated differently. For example, consider a security-enabled system with a type for users. The treatment of a user object depends on the user's privileges and possibly also on the past execution: Maybe we want to reduce the privileges when the user did not change the password for a while, or privileges are added or withdrawn at runtime in other ways.

Software design patterns [10] are another popular, more general example for dynamically varying the treatment of objects. Several design patterns define *roles* for objects, which can be assigned or removed at runtime, and the roles determine how an object is handled. However, while the pattern localizes the handling of object roles, the assignment of roles is usually scattered over multiple

source modules. As example, consider the *observer pattern*. To assign the role of *being observed* to a *subject*, an *observer* must be added to its observer list. The logic, which objects observe which subjects, is typically not implemented locally, but there are many different places in the code, where observers are added to subjects. Other similar examples are the *adapter*, *decorator*, or *proxy* patterns.

The need for influencing the handling of single objects, instead of all instances of a type, has already been discussed in 2003 by Rajan and Sullivan [19]. As a solution, they propose instance-level application of aspects. Similar solutions exist in AspectScheme [22], CaesarJ [1] or Steamloom [3]. However, all these approaches only support programmatic deployment of aspects on single instances; whereas our problem statement is that we need a way to modularly and declaratively define selection criteria for relevant instances.

More generally, we can say that objects have a history, how they were used, and we sometimes need to handle objects differently according to their past usage pattern. Often, the changes to an object's usage pattern are implicitly marked by events, e.g., passing an object from one client to another. We claim that to improve the modularity of source code, a declarative definition of the boundaries of relevant object-usage patterns is necessary. Furthermore, it must be possible to reify the set of objects that currently share a usage pattern to consider this information when handling an object.

Providing such a reification of object sets, requires modifying the source code by inserting bookkeeping code. Aspect-oriented programming (AOP) can be applied to separate this bookkeeping code from the business logic of the program. But in AOP, *pointcuts* select sets of so-called *join points*, which are points in time during the execution of the program. Current aspect-oriented languages do not support a *declarative specification* of a group of objects, which have been used in a similar way; instead an *imperative implementation*, typically following the same pattern, is required for collecting those objects.

A consequence of such an imperative solution is reduced readability and maintainability due to scattering, tangling and boilerplate code. Another issue is the lack of composition and checking mechanisms for the imperative bookkeeping. It is not possible to reuse the previously written code, which results in code that is hard to maintain and hinders software evolution. Also the warnings and errors do not indicate the proper context and relevant information to guide the programmer.

To offer better support for processing objects according to their usage history, we propose a new mechanism, called *instance pointcuts*, to select sets of objects based on the events in their execution history. Instance pointcuts are used to declare the beginning and the end of a relevant usage pattern as events. New instance pointcuts can be defined by reusing existing ones in two ways: first, by

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

PPPJ '14, September 23 - 26 2014, Cracow, Poland.
Copyright © 2014 ACM 978-1-4503-2926-2/14/09...\$15.00.
<http://dx.doi.org/10.1145/2647508.2647526>




Figure 1: Part of an online shop application

refining the expressions that define the relevant events and second, by *composing* two or more instance pointcuts using set operators.

In this paper we present a prototype of instance pointcuts as an extension to AspectJ [15] and explain its semantics by explaining our compiler, transforms instance pointcuts to plain AspectJ and advanced dispatching library calls.

We reuse the term *pointcut* for our concept, because it provides a declarative way of specifying crosscuts. Nevertheless, the instance pointcuts select *objects* whose usage crosscuts the program execution rather than *points (or regions) in time* [16] as traditional pointcut do. Therefore, our instance pointcuts cannot immediately be advised by AspectJ advice, although we offer the possibility to advise the points in time when the extent of instance pointcuts changes (cf. Section 3.4.2).

Our initial performance measurements show that maintaining object sets with our language extension is not slower than using plain Java. And at the same time, the declarative nature of instance pointcuts increases the modularity and extensibility of source code. It furthermore gives rise to several compile-time checks, which are not automatically possible with equivalent imperative code. Such checks are important to notify the developer when the instance pointcut set is guaranteed to be empty, incompatible types are used in compositions and refinements, etc. These checks help the developer to implement his concern correctly and achieve consistency in object sets representing the different usages of objects.

The rest of the paper is organized as follows, in Section 2 we present a small case study and explain our motivation for the proposed approach, and we formulate a problem statement and goals for our work. In Section 3, a detailed description of instance pointcuts and its various features are presented. Section 4 explains how instance pointcuts are compiled. We then present a discussion on the validation of our approach and outline possible checks in Section 5. We conclude by discussing related work and giving a summary of our approach.

2. Motivation

An important goal in programming language research is to enable modular extension of software, i.e., to facilitate that the implementation of new concerns can be added without modifying existing code. In general, the implementation of new concerns must interact with the already existing ones, and in object-oriented programming this is typically done through the exchange of objects. We have observed that often new concerns handle sets of objects from the original program to manipulate them or use them in a new way.

Objects can be grouped according to how they are used (where they are created, to which methods they are passed as argument, act as receiver or sender, etc.) and concerns of an application may be applicable only to objects used in a specific way. However, the relevant sets of objects may not be explicitly available in the original program. Thus, in this paper, we present a declarative approach for defining sets of objects with the same usage pattern by means of a dedicated language construct.

In Figure 1, we outline a part of the architecture of an online shop application. We use this scenario to give examples of grouping objects into sets according to how they are used, and how to use these sets in the implementation of concerns.

2.1 Example Architecture

In an online shop application, objects of the same type can exist, which nevertheless must be treated differently. In Figure 1 the static structure of a simplified online shop is shown. This structure shows part of the system from the Vendor and the OnlineShop's perspective. Vendors can submit different kinds of Discounts (not shown in the figure) to the ProductManager for the Products they are selling. Product is the root of the type hierarchy that represents different kinds of items that are sold in the online shop. Product is parent to classes such as BeautyProduct, SportProduct (not shown in the figure). Through the ProductManager, Discounts can be applied to and withdrawn from a Product. The OnlineShop has a user interface represented by the OnlineShopUI class, which is used to display information to the customers.

2.2 Extensions

A new feature is added to the online shop, which requires creating an alert when a product is applied a surprise discount. The list of surprise discounted products should be available to the user at any time. The surprise discounts are submitted by Vendors and they can be submitted or withdrawn at any time. In order to realize this extension in an OO-approach, we need to change several classes.

Listing 1 shows the modifications that must be made to the implementation of our online shop to accommodate for this new requirement. First, the class `ProductManager` should keep a set of `Products` to which a surprise discount is applied, as shown on line 3. This set is updated when a new discount of type `SurpriseDiscount` is submitted or withdrawn (lines 4–14). There should also be some changes in the `OnlineShop` class. A `createDiscountAlert` method should be added. Also the `displayList` method should be updated to include the `surpriseDiscount` list defined in the `ProductManager` class.

```

1  class ProductManager{
2      ...
3      Set<Product> surpriseDiscount = createSet();
4      public void submitDiscount(Product p, Discount d){
5          ...
6          if(d instanceof SurpriseDiscount){
7              surpriseDiscount.add(p);
8              OnlineShop.createDiscountAlert(p);
9          } }
10     public boolean withdrawDiscount(Product p, Discount d){
11         ...
12         if(d instanceof SurpriseDiscount)
13             surpriseDiscount.remove(p);
14     } }
15
16     class OnlineShop{
17         ...
18         public void createDiscountAlert(Product p){
19             //create surprise discount alert for p
20         }
21         public void displayList(String listType){
22             if(listType.equals("surprise"))
23                 INSTANCE.getUI().display(ProductManager.surpriseDiscount);
24             ...
25         } }

```

An aspect-oriented implementation can offer a better solution by encapsulating the concern in an aspect. Listing 2 shows a possible solution. The set of products, which are applied a surprise discount is kept in the aspect (line 2). The following two pointcuts `submit` and `withdraw` select the products to which a `SurpriseDiscount` is applied (lines 3–4). The corresponding advice declarations for

```

1 aspect SDiscount {
2     Set<Item> surpriseDiscount = createSet();
3     pointcut submit(Product p): call(* ↵
4         ProductManager.submitDiscount(..) && args(p, ↵
5             SurpriseDiscount);
6     pointcut withdraw(Product p): call(* ↵
7         ProductManager.withdrawDiscount(..) && args(p, ↵
8             SurpriseDiscount);
9     pointcut display(String listType): call(* OnlineShop.displayList(..) && ↵
10        args(listType);
11     after(Product p): submit(p){
12         surpriseDiscount.add(p);
13         OnlineShop.INSTANCE().createDiscountAlert(p);
14     }
15     after(Product p):withdraw(p){
16         surpriseDiscount.remove(p);
17     }
18     void around(String listType):display(listType){
19         if(listType.equals("surprise"))
20             OnlineShop.instance().getUI().display(surpriseDiscount);
21         proceed(listType);
22     }
23     public void OnlineShop.createDiscountAlert(Product p){
24         //create surprise discount alert for a product
25     }
26 }

```

Listing 2: An AspectJ implementation of discount alert concern

2.3 Discussion

AOP already helps to localize the concern of maintaining a collection of objects used in a specific way, and to add it without the need to modify existing code. However maintenance of the `surpriseDiscount` set requires the same boilerplate code as the OO solution does. Essentially the code selects `Product` objects based on the discount they are applied to and deselects them once they are rid of this discount. None of the solutions presented in this section offer declarative means to define such selection criteria. Furthermore, reusing the existing definitions of object-set reifications by refining or composing them is not conveniently supported at all; e.g., if we want to find the subset of `BeautyProducts` of the `surpriseDiscount` set, we have to iterate over it and check instance types to create a new set. Such imperative definitions are difficult or impossible to analyze by the compiler. For instance, it may be desirable to warn developers about instance pointcuts that provably will never match any object, e.g., because the selection events will never happen. With a more declarative notation, a compiler would be able to identify such situations.

2.4 Problem Statement

In this section we have demonstrated two things: First, the implementation of some concerns requires accessing groups of objects with similar usage history. Second, making such groups accessible to the program in a modular and re-usable way is not supported by current programming languages.

Since creating object sets according to execution events is a cross-cutting concern, we claim that a new programming technique in the style of aspect-oriented programming is required for modularizing concerns depending on object groups. Such a programming technique must satisfy the following needs:

- A declarative way of selecting/de-selecting objects according to the events they participate in should be provided.

- The selected objects should be accessible as a set.
- Changes to this set, i.e. adding or removing objects, should create a notification.
- For compatible element types, the sets should be composable to obtain new (managed) sets.

3. Instance Pointcuts

To support the requirements outlined in the previous section, we propose a new kind of pointcut for declaratively selecting objects based on their usage history. The beginning and the end of a relevant usage pattern is marked by events. *Instance pointcut* is a language construct that is used to reify and maintain a set of objects based on these events. It allows to modularize the object selection concern and makes it declarative.

In the remainder of this section, we will explain instance pointcuts in detail. Instance pointcuts can be implemented as an extension to arbitrary OO-based aspect-oriented languages. We have implemented a prototype as an extension to AspectJ. Therefore, the examples given throughout this paper are based on AspectJ.

A concrete instance pointcut definition consists of a left hand-side and a right-hand side (Figure 2, rule 1). At the left-hand side the pointcut's name and a type is declared. An instance pointcut does not declare pointcut parameters since it has the specific purpose of exposing one object from an event; it has a single implicit parameter called **instance** of the declared type.

At the right-hand side the instance pointcut expression selects the desired events from join points and then binds the exposed object (represented by the `instance` parameter) as a member of the instance pointcut's set.

$$\langle \textit{instance pointcut} \rangle ::= \text{'instance pointcut'} \quad \langle \textit{name} \rangle \quad '<'$$
$$\langle \textit{instance-type} \rangle '>' \text{' : ' } \langle \textit{ip-expr} \rangle (\text{'UNTIL'} \langle \textit{ip-expr} \rangle)?$$
$$\begin{array}{l} \langle ip\text{-}expr \rangle ::= \langle after\text{-}event \rangle '||' \langle before\text{-}event \rangle \\ \quad | \\ \quad \langle before\text{-}event \rangle '||' \langle after\text{-}event \rangle \\ \quad | \\ \quad \langle after\text{-}event \rangle \\ \quad | \\ \quad \langle before\text{-}event \rangle \end{array}$$
$$\langle \textit{after-event} \rangle ::= \textit{'after' '(\langle pointcut-expression \rangle)'}.$$
$$\langle \textit{before-event} \rangle ::= \textit{before} \textit{ } (\langle \textit{pointcut-expression} \rangle)$$

Figure 2: Grammar definition for instance pointcuts

3.1 Add/Remove Expressions

For each instance pointcut events must be selected at which objects are added to its set, otherwise the instance pointcut set would always be empty. Optionally, an instance pointcut can select events at which objects are removed from its set. The ‘add to set’ and ‘remove from set’ operations are implicitly performed when certain events specified in the corresponding *instance pointcut expression* (*ip-expr*, cf. Figure 2, rule 1) occur. The first expression is the mandatory *add expression*. After the **UNTIL** clause an optional expression called the *remove expression* can be defined.

In AspectJ join points mark *sites* of execution; a join point by itself does not define an event. Pointcut expressions select join points and pointcuts are used with advice specifications to select a particular event in that join point. As discussed by Masuhara et al. [16] such a region-in-time join-point model hinders re-use of pointcuts.

We combine pointcut expressions with advice specifiers and obtain *expression elements*. Each expression element contains a point-

cut expression, which matches a set of join points. Then, from these join points, according to the advice specifier, the before or after events are selected. Both add and remove expressions are composed of expression elements, which can be a *before element* or an *after element* (Figure 2, rule 3–4). In Figure 2 the second grammar rule depicts this. An instance pointcut expression (add/remove expression) contains at least one *expression element* and at most two.

Restricting the syntax to only allow the specification of a single before- or after event or a disjunction of both is without loss of generality: an event specification of the form `before(<expression-1>)|| before(<expression-2>)` is equivalent to the event specification `before(<expression-1> || <expression-2>)`. The same holds for the specification of after-events or the usage of conjunction instead of disjunction. Combining the specification of a before- and an after-event with `&&` is not supported as an event can never be “before” and “after” at the same time (see also [4]). We require users to specify instance-pointcuts in such a normalized form, i.e., with at most one before- and at most one after-event specification, so that we have a unique way of referencing the different parts of an instance pointcut definition, as will be further discussed in Section 3.3.

Rules 3 and 4 contain the *<pointcut-expression>* rule, which represents an AspectJ pointcut expression. In addition to the defined syntax, we also define two semantic rules for the rules *<after-event>* and *<before-event>*: First, it is mandatory to specify one binding predicate (*args*, *target* etc.) that binds the *instance* parameter, since it represents the object to be added or removed from the set. And second, there cannot be more than one such binding predicate.

The second restriction is defined because the binding predicate also constrains the applicability of the *<pointcut-expression>*; i.e., the pointcut expression only matches if the designated context value is type-compatible with the instance pointcut’s *<instance-type>*. If we would allow multiple binding predicates, the desired semantics would be unclear: should the whole pointcut expression match only if all bound values are type-compatible? Or should all type-compatible values be bound? If we encounter the need for binding multiple values in one *<after-event>* or *<before-event>*, then we will determine what is the desirable behavior and potentially extend our language to make this behavior configurable.

The binding predicates are extended to include the *returning* clause. The *returning* clause binds the value returned by a method or a constructor. In AspectJ the syntax is restricted and *returning* can only be used in an after advice, since the returned value is only available after a method finishes execution, this is also true for Instance Pointcuts. Although we do not have this restriction syntactically, we enforce that the *returning* clause is used only with the after event selector by means of a semantic check.

In an instance pointcut expression, it is only possible to *OR* a before event with an after event. The *before* clause selects the start of executing an operation (i.e., the start of a join point in AspectJ terminology) and the *after* clause selects the end of such an execution. For two operations that are executed sequentially, the end of the first and the start of the second operation are treated as two different events. Thus, the before and after clauses select from two disjoint groups of events and the conjunction of a before and an after clause will always be empty.

```
1 static instance pointcut surpriseDiscount<Product>:
2   after(call(* ProductManager.submitDiscount(..))
3     && args(instance, SurpriseDiscount))
4   UNTIL
5     after(call(* ProductManager.withdrawDiscount(..))
6       && args(instance, SurpriseDiscount));
```

Listing 3: A basic instance pointcut declaration with add and remove expressions

The instance pointcut in Listing 3 shows a basic example. The left-hand side of the instance pointcut indicates that the pointcut is

called *surpriseDiscount* and it is interested in selecting *Product* objects. On the right hand side, there are two expressions separated by the **UNTIL** keyword. The first one is the add expression. It selects the join-point marked by the method *submitDiscount* and from the context of this event it exposes the *Product* object with the *args* clause and binds it to the *instance* parameter. The second one is the remove expression and it selects the after event *withdrawDiscount* call and exposes the *Product* instance in the method arguments and binds it to the *instance*. This pointcut is the solution of the set maintenance problem presented in the motivation (Section 2).

It is not necessary that instances are added and removed in a specific order, like in a first-in-first-out or a last-in-first-out fashion. Rather, when either of the add or remove expressions matches a join point, the specified instance is bound and added to the instance pointcut, or removed respectively.

Note that instance pointcuts do not keep objects alive, as instance pointcuts are non-invasive constructs, which do not affect the program execution in any way. So even if the remove expression was not defined for the *surpriseDiscount* instance pointcut, when the *Product* instances are collected by the garbage collector, they are removed from the set.

3.2 Multisets

An instance pointcut reifies an object set as a *multiset*. A multiset allows multiple appearances of an object. Every contained object has a corresponding cardinality, which indicates its multiplicity in the set. Multiset makes sure that the maintained object sets are handled properly when an object enters and exits a relevant usage pattern recursively.

As example, consider again the instance pointcut shown in Listing 3 and replace the binding expression on lines 3 and 6 with `args(instance, Discount)`, such that products are selected with any applied discount. Assume that different discounts are submitted for the same product, then the product will be kept in the instance pointcut’s object set until the same amount of discounts is withdrawn from the product. If instance pointcuts only supported a set, then as soon as one discount is removed from a product, it would be removed and it would appear as if there are no more discounts on that product.

3.3 Refinement and Composition

Instance pointcuts can be referenced by other instance pointcuts. They can be refined in two ways and they can be composed together to create new instance pointcuts.

3.3.1 Referencing and Type Refinement

Instance pointcuts are referenced by their names. Optionally the reference can also take an additional statement for *type refinement*, which selects a subset of the instance pointcut that is of the specified type. Type refinements require that the refinement type is a subtype of the original instance type. For example the instance pointcut *surpriseDiscount* (Listing 3) can be refined as shown in Listing 4. The refinement expression selects the subset of *BeautyProduct* instances from the set of *Product* instances selected by the *surpriseDiscount* instance pointcut. The *surpriseDiscountBeauty* instance pointcut is defined using the result of this expression. Note that with this notation objects that are of a subtype of *BeautyProduct* will also be selected.

```
1 static instance pointcut surpriseDiscountBeauty<BeautyProduct>:
2   surpriseDiscount<BeautyProduct>;
```

Listing 4: A type refined pointcut

3.3.2 Instance Pointcut Expression Refinement

In Section 3.1 we have introduced the instance pointcut expression, which consists of two expressions (add and remove expressions).

We provide an expression refinement mechanism, which makes it possible to reuse parts of the existing instance pointcut expressions to create new ones. The add and remove expressions can be accessed individually to be extended by concatenating other primitive pointcuts, so-called *refinement expressions*, with boolean operators. We offer a naming convention to access parts of the instance pointcut expression with different granularity. Note that this syntax is only valid when used in the context of an expression refinement.

<ip-ref> When an instance pointcut is referenced directly then the refinement expression is composed with the pointcut expression in all of the before and after event selectors, in the add and remove expressions.

<ip-ref>.add{remove} This provides access at the expression level. The refinement expression is composed with the pointcut expressions in referenced add or remove expression's before and after event selectors.

<ip-ref>.add{remove}_after{before} This naming convention is used to access the pointcut expressions of the individual before and after event selectors and provides the finest granularity. In fact the other two access statements can be written in terms of this one, they just provide a short hand for the collective expression refinements.

It is possible to compose any primitive pointcut, except the binding predicates, with an expression. Although we chose not to restrict this aspect, some compositions will not be meaningful for selecting objects. For example composing an **execution** pointcut with an expression that already includes a **call** pointcut will result in a non-matching pointcut expression. This is further discussed in Section 5.1.

Let us explain the usage of the expression access by examples. The example shown in Listing 5 shows a reuse of the `surpriseDiscount`'s add and remove expressions to create a new instance pointcut. The newly created pointcut's expression can be accessed through the aforementioned naming conventions. In order to assign the expressions we use the `'.'` assignment operator, to provide a uniform syntax. It is not possible to assign a higher granularity statement to a lower granularity one; i.e. the following is illegal `add_before : <ip-ref>.add`.

```
1 static instance pointcut surpriseDiscountOver50<Product>:
2   add: surpriseDiscount.add && if(instance.getPrice() > 50) UNTIL
3   remove: surpriseDiscount.remove;
```

Listing 5: Expression refinement of `surpriseDiscount`(Listing 3) instance pointcut

The `if` pointcut in Listing 5 is appended to the `add` expression of `surpriseDiscount` (Listing 3). The effect of this composition is as follows; the `if` pointcut will be appended to all of the pointcut expressions contained in the **after** and **before** event selectors. Since the `surpriseDiscount` pointcut only has one after event in its add expression, the resulting add expression is equivalent to:

```
after(call(* ProductManager.submitDiscount(...)) &&
args(instance, SurpriseDiscount) && if(instance.getPrice() > 50))
```

Expression refinements can also be used for more precise type refinements. Revisiting the example given in Subsection 3.3.1, the `surpriseDiscountBeauty` (Listing 4) instance pointcut can be constructed to include instances with the *exact type* `BeautyProduct` (Listing 6). The effect is different from type refinement since `surpriseDiscountOnlyBeauty` does not include subtypes of the class `BeautyProduct`.

```
1 static instance pointcut surpriseDiscountOnlyBeauty<BeautyProduct>:
2   surpriseDiscount && if(instance.getClass().equals(BeautyProduct.class));
```

Listing 6: Type refinement by expression refinement

3.3.3 Instance Pointcut Composition

Instance pointcuts reify sets, for this reason we facilitate the composition in terms of set operations: *intersection and union*. In Figure 3, an extended version of the grammar definition is shown. The composition of two instance pointcuts creates a *composite* instance pointcut. Different from regular instance pointcuts, composite ones are declared with the keyword **composite** and they do not have instance pointcut expressions. Instead they monitor the component instance pointcuts' set change operations and update their own set accordingly. In order to declare a set intersection the keyword **inter** and to declare a set union the keyword **union** is used. Throughout the text we will use the mathematical symbols for these operations, \cap as intersection and \cup as union. Since composite instance pointcuts do not have an instance pointcut expression they cannot be used in expression refinement. However they can be type-refined; the result of the type refinement of a composite instance pointcut is also a composite instance pointcut and must be declared as such.

```
<instance-pointcut> ::= 'composite instance pointcut'
  <name> ('<' <instance-type> '>')? ':' ...
  | <comp-expr>

<comp-expr> ::= <comp-expr> 'inter' <comp-expr-t>
  | <comp-expr-t>

<comp-expr-t> ::= <comp-expr-t> 'union' <comp-expr-f>
  | <comp-expr-f>

<comp-expr-f> ::= <ip-ref>
  | '(' <comp-expr> ') '


<ip-ref> ::= <name>
  | <name> ('<' <refined-instance-type> '>')?
```

Figure 3: Syntax for instance pointcut composition

The type of a composite instance pointcut must be assignment compatible to the types of the component instance pointcuts. It is also possible to leave out the type declaration and let the compiler assign the type. For a composition of two instance pointcuts, the type of the composite one can be determined depending on the relation of the types of the component instance pointcuts. For illustration of this type inference, consider the type hierarchy in Figure 4a: *R* is the root of the hierarchy with the direct children *A* and *B* (i.e., these types are siblings); *C* is a child of *B*. Table 4b shows four distinct cases: Either the type of one of the instance pointcuts is a super type of the other one's type (second row), or both types are unrelated (third row); and the composition can either be \cap (third column) or \cup (fourth column).

When composing two instance pointcuts with types from the same hierarchy, the type of the composition is the more specific type (*C* in the example) for an \cap composition and the more general type (*B*) for an \cup composition. When composing two instance pointcuts with sibling types, for the \cap operation the resulting composition cannot select any types since the types *A* and *B* cannot have a common instance. The \cup operation will again select a mix of instances of type *A* and *B*, thus composed instance pointcut must have the common super type, *R* in the example.

Because instance pointcuts are reified as multisets, these operations are different from the regular set operations. The definition of the intersection and union operations for multiset is given in the next definition.



(a) A simple type hierarchy

pc1 type	pc2 type	\cap	\cup
B	C	C	B
B	A	\emptyset	R

(b) Instance pointcut compositions and effect on the captured instance type

Figure 4: An example to illustrate composition's effect on types

DEFINITION 1. Assume (X, f) and (Y, g) are multisets, where X, Y represents the elements and f, g represents a function, which maps each element to a cardinal number.

The **intersection** of these sets is defined as (V, h) where,

$$V = X \cap Y$$

and $\forall v \in V$ the multiplicity of v is defined as

$$h(v) = \min(f(v), g(v))$$

The **union** of these sets is defined as (Z, i) where,

$$Z = X \cup Y$$

and $\forall z \in Z$ the multiplicity of z is defined as

$$i(z) = \max(f(z), g(z))$$

3.4 Using Instance Pointcuts

Up to now we have explained the syntax and semantics for definitions of instance pointcuts. In this section we will explain how to use instance pointcuts in the context of an AO language, namely, AspectJ. As example, throughout this section, we will use the instance pointcut defined in Listing 7, which maintains a set of Products that are currently out of stock. Instance pointcuts are static members of classes and can have any visibility modifier. Thus, all modules, aspects as well as classes, that can see an instance pointcut can use it in the ways described below.

```

1  static instance pointcut outOfStock<Product>:
2      after(call(* Product.outOfStock(...)) &&
3          target(instance))
4  UNTIL
5      after(call(* Vendor.stock(...))
6          && args(instance));

```

Listing 7: An instance pointcut for out of stock products

3.4.1 Set Access

Instance pointcuts reify a set of objects, which are used in a similar way, and this set can be accessed through a static method, which has the same name as the instance pointcut identifier. Only the *get* methods of the collection interface can be used to retrieve objects from the set. We ensure this by returning an *UnmodifiableSet* from the set access methods. In Listing 8 the *outOfStock()* method (line 3) returns the set of Products that are currently out of stock. Write methods, which modify the contents of the set, are not allowed since they create data inconsistencies like adding an object, which

is not used in the same way as the ones selected by the instance pointcut. The usage of write methods may result in concurrent modification exceptions.

```

1  public static double calculateDamages() {
2      double damage = 0;
3      for(Product p: MyAspect.outOfStock())
4          damage = damage + p.getPrice();
5      return damage;
6  }

```

Listing 8: Calculate a damage estimate for out of stock products

3.4.2 Set Monitoring

An instance pointcut definition defines two set change events, an add event and a remove event. In order to select the join points of these events, every instance pointcut definition automatically has two implicit regular pointcuts. These implicit pointcuts have the following naming conventions, $\langle name \rangle_instanceAdded$, $\langle name \rangle_instanceRemoved$, where $\langle name \rangle$ is the name of the instance pointcut. In Listing 9, a before advice using the *outOfStock_instanceAdded* pointcut is shown. When a product is marked out of stock and it is added to the set, a notification is sent to the related Vendor indicating that the product is out of stock.

```

1  before(Product p): outOfStock_instanceAdded(p) {
2      OnlineShop.notifyVendor(p.getVendor, STOCK_MSG);
3  }

```

Listing 9: Set monitoring pointcut used to notify vendors

4. Compilation of Instance Pointcuts

A goal for our compiler implementation is to support modular compilation. This means to compile an aspect with instance pointcuts that refer to instance pointcuts defined in other aspects, it must be sufficient to know their declaration (i.e., the name and type); it should not be necessary for the compiler to know the actual expression or the referenced instance pointcuts.

We have implemented the instance pointcut language using code transformation employing two tools. First, the parser of our language and the code generation templates are implemented with the EMFText¹ language workbench. For this purpose, we have defined the AspectJ grammar by using JaMoPP² [14] as the foundation and extended it with the grammar for instance pointcuts, which was presented interspersed with the previous section.

Second, the generated code uses the ALIA4J³ [6] framework for so-called advanced dispatching language implementations. The term advanced dispatching refers to late-binding mechanisms including, e.g., predicate dispatching and pointcut-advice mechanisms. At its core, ALIA4J contains a meta-model of advanced dispatching declarations, called *LIAM*, in which AspectJ pointcut and advice, as well as instance pointcuts can be expressed.

We use ALIA4J to realize the crosscutting behavior of our language instead of AspectJ because the way AspectJ handles binding of values and restricting their types in pointcuts would prohibit a modular compilation of instance pointcuts. While in instance pointcuts the value binding is uniformly expressed in a pointcut expression, in AspectJ binding the result value must be specified in the advice definition (via the **after returning** keyword) and all other values are bound in pointcut expressions. Therefore, AspectJ code generated for an instance pointcut expression would have to depend on

¹EMFText, see <http://www.emftext.org/>

²JaMoPP: Java Model Parser and Printer, see <http://jamopp.inf.tu-dresden.de>

³The Advanced-dispatching Language Implementation Architecture for Java. See <http://www.alia4j.org/alia4j/>.

which value is bound; this means that the code generation for a derived instance pointcut would also depend on the binding predicate (an implementation detail) of the referenced one. It is not possible to work around this, using AspectJ's reflective `thisJoinPoint` keyword, as it does not expose the result value at all. Another, similar limitation is that AspectJ does not allow to narrow down the type restriction for the bound value of a referred pointcut. Thus, in order to be able to transform an instance pointcuts with type refinement to AspectJ, it is necessary to know the definitions of the referenced instance pointcut and inline its definition.

Our compiler generates different code depending on whether the instance pointcut is a composite one or not and whether it is a refinement of an instance pointcut or not. Common to all cases is the code for managing the data of the instance pointcut. Listing 10 exemplary shows that code; the variables `#{Type}` and `#{ipc}` stand for the instance pointcut's type and name, respectively. The natural text written in the comments provides a description of the code for which it stands.

First, to store the instances currently selected by an instance pointcut as a multiset, a `WeakHashMap` is defined (cf. line 1); the keys of the map are the selected objects and the mapped value is the cardinality. We use weak references to avoid keeping objects alive, which are not reachable from the base application anymore. The generated method `#{ipc}` returns all objects, which are currently mapped (cf. lines 2–4).

Furthermore, methods are generated to access, increase or decrease the counter of selected objects; if an object does not have an associated counter yet or the counter reached zero, the object is added to or removed from the map, respectively (cf. lines 5–15). After having performed their operations, `#{ipc}_addInstance` and `#{ipc}_removeInstance` methods invoke an empty method, passing the added or removed object. We generate a public, named pointcut selecting these calls, exposing the respective events (cf. lines 18 and 19).

```

1 private static WeakHashMap<#{Type}, Integer> #{ipc}_data = new
  WeakHashMap<#{Type}, Integer>();
2 public static Set<#{Type}> #{ipc}() {
3     return Collections.unmodifiableSet(#{ipc}_data.keySet());
4 }
5 public static void #{ipc}_addInstance(#{Type} instance) {
6     //increase counter associated with instance by the #{ipc}_data map
7     #{ipc}_instanceAdded(instance);
8 }
9 public static void #{ipc}_removeInstance(#{Type} instance) {
10    //decrease counter associated with instance by the #{ipc}_data map
11    //if the counter reaches 0, remove instance from the map
12    #{ipc}_instanceRemoved(instance);
13 }
14 public static int #{ipc}_cardinality(#{Type} o) {...}
15 private static void #{ipc}_setCardinality(#{Type} o, int c){...}
16 private static void #{ipc}_instanceAdded(#{Type} instance) {}
17 private static void #{ipc}_instanceRemoved(#{Type} instance) {}
18 public pointcut #{ipc}_instanceAdded(#{Type} instance) : call(private
  static void Aspect.#{ipc}_instanceAdded(#{Type})) &&
  args(instance);
19 public pointcut #{ipc}_instanceRemoved(#{Type} instance) :
  call(private static void
  Aspect.#{ipc}_instanceRemoved(#{Type})) && args(instance);

```

Listing 10: Template of generated code for instance set management.

Next, these bookkeeping methods have to be executed at events corresponding to the instance pointcut definitions. Below, we elaborate on the code generation for instance pointcuts defined in the different possible ways.

4.1 Non-Composite Instance Pointcuts

A non-composite instance pointcut, generally consists of four underlying pointcut definitions: specifying join points (1) *before* or




Figure 5: Meta-model of a Specialization in ALIA4J.

(2) *after* which an instance is to be *added* to the selected instances; and specifying join points (3) *before* or (4) *after* which an instance is to be *removed*. For each pointcut definition, we generate a method that creates a corresponding LIAM model; the methods are called `#{ipc}_add_before`, `#{ipc}_add_after`, `#{ipc}_remove_before`, and `#{ipc}_remove_after`.

In LIAM, a *Specialization* can represent a partial AspectJ pointcut and a full pointcut expression can be represented as the disjunction of a set of *Specializations* (discussed in detail elsewhere [5]). Figure 5 shows the meta-model for a *Specialization* in ALIA4J consisting of three parts. A *Pattern* specifies syntactic and lexical properties of matched join point shadows. The *Predicate* and *Atomic Predicate* entities model conditions on which the dynamic state of pointcut designators depend. The *Context* entities model access to values like the called object or argument values. Contexts which are directly referred to by the *Specialization* are exposed to associated advice (i.e., they represent binding predicates).

Depending on the definition of the instance pointcut, LIAM models of the underlying pointcuts have to be created in different ways. All four underlying pointcuts are optional; a missing pointcut can be represented as an empty set of *Specializations* in LIAM.

Plain Instance Pointcuts For pointcut expressions that are directly provided, we use a library function provided by ALIA4J, which takes a String containing an AspectJ pointcut as input. We have extended this library to also accept the *returning* pointcut designator.

```

1 static instance pointcut #{ipc}<#{Type}>:
2 before(#{pc_add_before}) after(#{pc_add_after}) UNTIL
3 before(#{pc_remove_before}) after(#{pc_remove_after});

```

Listing 11: Example of a plain instance pointcut

For the example instance pointcut presented in listing 11, we show the code generated for the method creating the LIAM model for the `add_before` pointcut in listing 12; the other methods are generated analogously. Line 4 shows the transformation of an AspectJ pointcut into a set of *Specializations* in the LIAM meta-model by passing the pointcut as a String—represented by `#{pc_add_before}` in listings 11 and 12—to the above mentioned library function.

```

1 private static Set<Specialization> #{ipc}_add_before;
2 public static Set<Specialization> #{ipc}_add_before() {
3     if (#{ipc}_add_before == null) {
4         #{ipc}_add_before = Util.toSpecializations("#{pc_add_before}",
5             #{Type});
6     }
7     return #{ipc}_add_before;

```

Listing 12: Template for creating the LIAM model for the `add_before` expression

Type Refinement An instance pointcut can also be defined by referring to another instance pointcut whereby a type restriction for the selected instances can be defined. A template for an instance pointcut defined as such is as follows:

```
static instance pointcut ${ipc} < ${Type} >: ipc1 < ${Type} >;
```

For such an instance pointcut also methods are created to produce LIAM models for the four underlying pointcuts (see the template in listing 13). These methods first invoke the corresponding methods of the referenced instance pointcut (cf. line 3). Second, the type restriction is added to the Predicates of the retrieved Specializations (cf. line 4) with the method `addTypeConstraint`.

```
1 public static Set<Specialization> ${ipc}_add_before() {
2   ...
3   Set<Specialization> ipRef = ipc1_add_before();
4   ${ipc}_add_before = Util.addTypeConstraint(ipRef, ${Type});
5   ...
6 }
```

Listing 13: Template for creating the LIAM model for the type-refined instance pointcut

Expression Refinement When defining a new instance pointcut through expression refinement, for each of the four underlying pointcut expressions, a plain pointcut expression can be *anded* or *ored* with the underlying pointcut expression of the referenced instance pointcut. As explained in the previous section, refinement pointcut expressions must not include a binding predicate. The referred instance pointcut expression already has a binding predicate, which is carried over to be used as the binding predicate for the newly composed pointcut expression. The template for the generated method for creating the add/before LIAM model is shown in listing 14. It assumes that the instance pointcut is named `${ipc}` and it refines another instance pointcut `${ipc1}` by *anding* the pointcut expression `${ipc1}_add_before` to the add/before underlying pointcut. If the pointcuts are *ored*, correspondingly the method `orSpecializations` is used in line 5.

These utility methods are provided as runtime library for the instance pointcuts. The method `andSpecializations` forms the conjunction of the Predicates and Patterns of the passed Specialization sets. If `ipRef` is empty, the conjunction is also empty and an empty set is returned. Otherwise, the Context declared by the Specializations `ipRef` is copied to the ones newly created by the `andSpecializations` method. The method `orSpecializations` is implemented similarly. But it forms the disjunction of Predicates and Patterns and if `ipRef` is empty, an exception is raised.

```
1 public static Set<Specialization> ${ipc}_add_before() {
2   ...
3   Set<Specialization> plainIPEXpr = 
4       Util.toSpecializations("${pc}_add_before", ${Type});
5   Set<Specialization> ipRef = ipc1_add_before();
6   ${ipc}_add_before = Util.andSpecializations(ipRef, plainIPEXpr);
7   ...
8 }
```

Listing 14: Generated code for creating the LIAM model for the add/before pointcut of the instance pointcut created with expression refinement

Deployment In each of the above cases, the created LIAM models of the pointcuts must be associated with advice invoking the add or remove method for the instance pointcut. In a LIAM model this is achieved by defining an Attachment, which roughly corresponds to a pointcut-advice pair. An Attachment refers to a set of Specializations, to an *Action*, which specifies the advice functionality, and to a *Schedule Information*, which models the time relative to a join point (cf. line 6).

Listing 15 shows the generated code for creating and deploying the bookkeeping Attachments. The first Attachment uses the set of Specializations returned by the `${ipc}_add_before` method (cf. line 4) and specifies the `${ipc}_addInstance` method as action to execute

at the selected join points (cf. line 5). As relative execution time, the Attachment uses a “SystemScheduleInfo”; this is provided by ALIA4J for Attachments performing maintenance whose action should be performed before or after all user actions at a join point, such that all user actions observe the same state of the maintained data. Thus, when reaching a selected join point the instance is added to the instance pointcut’s multiset before any other action can access its current content. The other Attachments are created analogously. In the end, all Attachments are deployed through the ALIA4J System (cf. line 2).

```
1 public static void ${ipc}_deploy() {
2   org.alia4j.fial.System.deploy(
3     new Attachment(
4       ${ipc}_add_before(),
5       createStaticAction(void.class, ${Aspect}.class, 
6         "${ipc}_addInstance", new Class[] { ${Type}.class })
7       SystemScheduleInfo.BEFORE_FARTHEST),
8     //Create Attachments for the other three parts analogously.
9     //For the "after" parts, use SystemScheduleInfo.AFTER_FARTHEST.
10    //For the "remove" parts, specify method ${ipc}_removeInstance.
11  );
12 }
```

Listing 15: Deployment of the bookkeeping for an instance pointcut.

4.2 Composite Instance Pointcuts

Composite instance pointcuts require a different compilation strategy because they do not have the four underlying pointcut expressions. The data of a composite instance pointcut changes when the data of one of its referenced instance pointcuts is updated. The corresponding events happen during the execution of the generated methods `${ipc}_addInstance` and `${ipc}_removeInstance`. Therefore, a different mechanism is needed than for the non-composite instance pointcuts, which depend on user events.

Intersection and Union When an instance pointcut `${ipc}` is composed by forming the union or intersection of other instance pointcuts (`${ipcX}`), the content of the maintained multiset potentially changes whenever an instance is added to or removed from one of the referenced instance pointcuts—events already exposed through `${ipcX}_instanceAdded` or `${ipcX}_instanceRemoved` pointcuts. For the maintenance of a composite instance pointcut a method is generated, which reacts to the join points matching the disjunction of all these pointcuts. The argument of this method is the instance exposed by these pointcuts, i.e., the instance that has either been added to or removed from a reference instance pointcut. When the maintenance method is invoked, we know the cardinality of this instance potentially changes in the multiset of the instance pointcut `${ipc}`. The cardinality of other instances cannot change. The generated method, therefore, re-calculates the cardinality of the affected instance and changes its value in `${ipc}_data`.

To generate appropriate code, the compiler first builds a binary expression tree for the composition expression. Next, it traverses this tree and generates different code for the cases that the visited node is an instance pointcut reference, or an `inter` or `union` operator. For an instance pointcut reference, code is generated that retrieves the cardinality of the instance in the multiset of the referenced instance pointcut. For an `inter` and `union` operator, code is generated that calculates the minimum and maximum, respectively, of both expressions. Finally, the cardinality in `${ipc}_data` is updated.

As example, listing 16 shows the generated code for a composite instance pointcut with the set expression `${ipc1} union ${ipc2} inter ${ipc3}`. Besides, the generated code remembers the old cardinality; when the cardinality changes from 0 to > 0 or vice versa, the generated method invokes the method `${ipc}_instanceAdded` `${ipc}_instanceRemoved`, respectively.


```

1 public static void ${ipc}_update(Object o) {
2     int oldCardinality = ${ipc}_cardinality(o);
3     int newCardinality =
4         Math.min(
5             Math.max(
6                 ${ipc1}_cardinality(o),
7                 ${ipc2}_cardinality(o)),
8                 ${ipc3}_cardinality(o));
9     ${ipc}.setCardinality(o, newCardinality);
10    if (oldCardinality == 0 and newCardinality > 0)
11        ${ipc}_instanceAdded(o);
12    else if (oldCardinality > 0 and newCardinality == 0)
13        ${ipc}_instanceRemoved(o);
14 }

```

Listing 16: The update method generated from a composition expression

As in the case of non-composite instance pointcuts, a LIAM Attachment is generated and deployed, which associates the Specializations corresponding to the pointcuts with the generated method.

Type Refinement for Composite Instance Pointcuts Instance pointcuts which are defined by means of type refined composite instance pointcuts, are treated similar to the case above. A method is generated which is executed when the referenced instance pointcut changes. The method checks whether the type of the added or removed object is assignment compatible with the type restriction. If this is the case, the same operation (adding or removing the instance) is performed on the multiset of the refining instance pointcut.

4.3 Compiling Plain AspectJ constructs

To ensure consistent ordering between AspectJ advice and our implementation of instance pointcuts (i.e., that our bookkeeping advice are executed before user advice), the AspectJ pointcut-advice definitions must be processed by ALIA4J. This is possible because ALIA4J can integrate with the standard AspectJ tooling. Using command line arguments the AspectJ compiler can be instructed to omit the weaving phase. The advice bodies are converted to methods and pointcut expressions are attached to them using Java's annotations, which are read by the ALIA4J-AspectJ integration and transformed into Attachments at program start-up. The code generated by our compiler consists of the above explained methods, as well as plain AspectJ definitions. When compiling this code with the mentioned command line options, the regular AspectJ pointcut-advice and the behavior of the instance pointcuts are both executed by ALIA4J, thus ensuring a consistent execution order.

5. Validation

5.1 Enabled Checks

The instance pointcuts approach satisfies the goals we have stated at the beginning of Section 2.4 and in Section 4 we have shown that instance pointcuts can be compiled modularly. Our approach provides:

- A concise syntax, which is used to generate the necessary bookkeeping code,
- additional features to declaratively create refined sets, reusing already created instance pointcuts, and
- a composition mechanism, which uses set operations, that allows modular definition of instance pointcuts.

Let us revisit the problem we have identified in Section 2 and provide an implementation that uses instance pointcuts shown in Listing 17 (this listing only includes the changed code, the rest is identical to Listing 2). The instance pointcut `surpriseDiscount` successfully encapsulates the bookkeeping concern, which reduces the

concern specific lines of code to a third of the AspectJ solution. The discount alert advice now uses the `surpriseDiscount_instanceAdded` set monitoring pointcut (lines 4–6). Also, in order to access the instance pointcut set, the code now uses the set access method as in line 10. Introducing new concerns to this code, like refining the `surpriseDiscount` instance pointcut to select objects with prices that are higher than 50\$ (Listing 5) requires very little amount of code.

```

1 aspect SDiscount{
2     static instance pointcut surpriseDiscount<Product>: after(call(* ✓
3         ProductManager.submitDiscount(...) && args(instance, ✓
4             SurpriseDiscount))
5         UNTIL after(call(* ProductManager.withdrawDiscount(...) && ✓
6             args(instance, SurpriseDiscount));
7     after(Product p):surpriseDiscount_instanceAdded(p){
8         OnlineShop.INSTANCE().createDiscountAlert(p);
9     }
10    void around(String listType):display(listType){
11        ...
12        OnlineShop.instance().getUI()
13        .display(SDiscount.surpriseDiscount());
14    }...
15 }

```

Listing 17: The instance pointcut implementation of the discount alert concern

A declarative syntax such as that of instance pointcuts, generally allows performing various checks, generate well-placed error or warning markers and informative warning/error messages. In the following we discuss checks we deem useful and possible to implement based on our language. Proving their feasibility by implementing the checks is subject to future work.

Non-existent/Incompatible types If the type declared by the instance pointcut does not exist, this is a compile error. Instance pointcuts provide an additional check during type refinement; it is a compile error if the refinement type is not a subtype of the referenced pointcut's declared type. In the Java solution the error marker is placed at the line of the `instanceof` check, without giving any context to why the `instanceof` check is performed. The instance pointcut error marker is placed at line of the refinement, with an elaborate error message explaining the type mistake.

Type compatibility is also an issue while composing instance pointcuts. The typing effects of the composition is previously mentioned in Section 3.3.3. If in a composite instance pointcut declaration, the instance type is explicit, then the type compatibility between the composite and the component instance pointcuts' should be checked. The compiler determines the appropriate type for every composite instance pointcut, whether the instance type is explicit or not. The computed type is then compared to the declared type; the declared type must be the same or a super type of the computed type. If there's an incompatibility, we put an error marker at the line where the type was declared, indicating the composed type is not compatible with the declared type of the composite instance pointcut.

Empty Sets The add expression of an instance pointcut is responsible for populating the instance pointcut set. If the pointcut expressions defined in the add expressions do not match any join-points then it is guaranteed that the instance pointcut set will always be empty. This case is displayed as a compile-time warning, which indicates that no objects will be selected. During the expression refinement, composition of the new pointcuts and the referenced one may result in non-matching pointcut expressions. Currently we do not perform any control-flow analysis to check for empty sets. Such a check will help us to identify cases where even if the add expression matches, no objects are added to the set.

Double selection Instance pointcut syntax allows selecting the same join-point in both before and after events in the same expression. In Listing 18 such a case is illustrated. The `Product` object is

added before it is applied a discount, and once again after the discount is applied. This will increment the cardinality of the same object twice. The analogous case exists for the remove expressions. The object is removed from the set *twice*. This behaviour is consistent with an instance pointcut's regular behaviour, but we still raise a separate warning for the add and remove cases to notify the developer, in case of a copy-paste error.

```
1 static instance pointcut surpriseDiscountDouble<Product>:
2   before(call(* ProductManager.submitDiscount(..))
3     && args(instance, SurpriseDiscount))
4   ||
5   after(call(* ProductManager.submitDiscount(..))
6     && args(instance, SurpriseDiscount))
```

Listing 18: Adding the same object before and after the same joinpoint

Expression Refinement Checks During expression refinement there is a special case when the refined instance pointcut is referencing a non-existent event selector and the `||` boolean operator is used. Assume that `ipc1` only has an `add_before` expression. While refining this instance pointcut the developer mistakenly writes the following:

```
1 ipc2<T>: add_before:ipc1.add_before
2   add_after: ipc1.add_after || call(..);
```

This definition refers to the non-existent `add_after` expression. This is a compile error since, while the `call` pointcut is selecting joinpoints, no objects are bound. This is illegal to have in an instance pointcut expression. Note that if the operation was `&&` there would be no compile error since, the `add_after` expression of `ipc2` would simply be empty and `ipc2`'s `add` expression would only comprise of `ipc1`'s `add_before` event selector.

5.2 Real-World Example

The Git Repository Hosting site called github offers a mobile application to access its services via Android smart phones. We have performed a preliminary study on this application to show the benefits of using instance pointcuts in a real-life application. An Android application consists of Activity objects which are application components users can interact with to perform a task. Each activity is further divided into Fragments, which represents a portion of that Activity. Fragments have a number of states which mark the beginning and the end of their lifecycle. The users of github can be associated with certain organisations, which gives them access to repositories which are not publicly available but only available through organisation permissions.

The classes which were affected by the addition of this feature can be seen on commit with ID 044262d on the git repository of the project.⁴ This particular commit includes changes to `HomeActivity`, `RepositoryListFragment`, `MembersFragment` and `UserNewsFragment`. It also adds two new interfaces `OrganizationSelectionListener` and `OrganizationSelectionProvider`. All of these classes except `RepositoryListFragment` belong to the same package. All of the listed fragments start their lifecycle with a call to the `onActivityCreated` and they end their lifecycle with a call to the `onDetach`. Fragments register themselves as a `OrganizationSelectionListener` when they are created to the `HomeActivity`. In the implementation of these methods we see the same bookkeeping code repeated in the listing below. The purpose of this code is to listen to the changes of the selected organisation for the active user and update the related UI fragments accordingly.

```
1 [RepositoryListFragment,MembersFragment,UserNewsFragment] extends <
   Fragment implements OrganizationSelectionListener {
```

⁴ See <https://github.com/github/android/commit/044262d>.

```
2 ...
3 @Override //Fragment
4 public void onDetach() {
5     OrganizationSelectionProvider selectionProvider =
6         (OrganizationSelectionProvider) getActivity();
7     if (selectionProvider != null) selectionProvider.removeListener(this);
8     super.onDetach(); }
9 @Override //Fragment
10 public void onActivityCreated(Bundle savedInstanceState) {
11     org = ((OrganizationSelectionProvider) getActivity()).addListener(this);
12     //rest of the implementation is different in each class ...
13 }
14 @Override //OrganizationSelectionListener
15 public void onOrganizationSelected(final User organization) {...} }
```

The `HomeActivity` class implements the `OrganizationSelectionProvider` interface which offers the `addListener(OrganizationSelectionListener)` and `removeListener(OrganizationSelectionListener)` methods. The added listeners are kept in the `orgSelectionListeners` set. When a user selects an organisation, each listener is notified in a for loop by calling their `onOrganizationSelected` method. From this implementation we can see that the bookkeeping of these listeners is scattered among 3 classes, a total of 6 methods. In order to modularise this concern, we have replaced this listener add/remove implementation with an instance pointcut which is shown in the listing below.

```
1 public aspect OrganizationSelectionProviderAspect{
2   static instance pointcut <
3     orgSelectionListeners<OrganizationSelectionListener>:
4     before(call(* RepositoryListFragment.onActivityCreated(..))
5       call(* MembersFragment.onActivityCreated(..))
6       call(* UserNewsFragment.onActivityCreated(..) && this(instance)) <
7     UNTIL
8     before(call(* RepositoryListFragment.onDetach())
9       call(* MembersFragment.onDetach()))
10    call(* UserNewsFragment.onDetach()) && this(instance));
11 }
```

This has resulted in the following code quality improvements in the affected classes.

The overridden `onDetach` method is removed from all fragment classes, since its only purpose was to remove the listeners from the provider. This is now managed by the `orgSelectionListener` instance pointcut. This is a quality improvement since we were able to eliminate redundant code from these classes. This has reduced a total of 21 lines of code from three classes.

- The first line of the `onActivityCreated` method was removed, reducing 3 lines in total.
- We have removed the `OrganizationSelectionProvider` interface, the implementation of its methods and the `orgSelectionListeners` set from `HomeActivity` class, reducing in total 11 lines of code.
- The removal of the `OrganizationSelectionProvider` interface also makes the dependency to this class obsolete and makes it possible to remove this dependency from the affected Fragment classes. The removal of this interface also simplifies the type-hierarchy.

We were able to localise the bookkeeping of the listener in a single module resulting in more maintainable code; instead of altering three different classes the developer can write a single instance pointcut. Besides modularising the implementation of this concern, the solution with instance pointcuts is also more concise. The 35 lines of the scattered implementation were replaced by 16 lines for an aspect with an instance pointcut. This reduced the size of the implementation of this concern by more than 50%.

5.3 Performance

We have used the Caliper [12] micro-benchmarking framework for measuring the execution time of our language extension. In this experiment, we measured the performance of add and remove

SCENARIO, BENCHMARKSPEC, PARAMETERS, SIZE	SCENARIO, BENCHMARKSPEC, METHODNAME	RUNTIME (NS)
1	testAddUniqueIpc	35,458
1	testAddUniqueJava	35,746
10	testAddUniqueIpc	361,996
10	testAddUniqueJava	358,938
100	testAddUniqueIpc	3,777,122
100	testAddUniqueJava	3,793,334
1000	testAddUniqueIpc	38,376,621
1000	testAddUniqueJava	37,949,489
10000	testAddUniqueIpc	358,668,568
10000	testAddUniqueJava	363,596,696

Figure 6: Benchmark results for adding unique objects

SCENARIO, BENCHMARKSPEC, PARAMETERS, SIZE	SCENARIO, BENCHMARKSPEC, METHODNAME	RUNTIME (NS)
1	testAddSameIpc	35,864
1	testAddSameJava	36,827
10	testAddSameIpc	369,966
10	testAddSameJava	368,380
100	testAddSameIpc	3,642,220
100	testAddSameJava	3,597,256
1000	testAddSameIpc	36,343,886
1000	testAddSameJava	36,719,268
10000	testAddSameIpc	361,327,587
10000	testAddSameJava	356,892,757

Figure 7: Benchmark results for adding the same object

operations of instance pointcuts, once with our implementation and once with semantically equivalent plain Java code. The two scenarios we have measured are as follows:

- *Adding/removing unique objects* to the instance pointcut set will result in a growing data structure with objects of cardinality value one, similarly removing these objects will shrink the data structure.
- *Adding/removing the same object* more than once to the instance pointcut set will result in a set with a single element and a cardinality greater than one. The remove operation will only reduce this cardinality, until the cardinality becomes zero. Then the object is removed from the set.

We have performed our experiments on a laptop running MacOSX v10.9.1 with 2.4 Ghz Intel Core i7 processor and 8GB of RAM with the Java version JRE 1.7.0u40 64-bit. We have used the default setting of Caliper and instructed the framework to record 10 measurements for each benchmark.

Figure 6 shows the results for adding unique objects to the instance pointcut. And Figure 7 shows the results for adding the same object. We show the results for adding objects only and the results for removing objects are equivalent.

In these results, the average of the plain Java case is slightly faster than using instance pointcuts. Nevertheless, the confidence intervals overlap and thus it cannot be said that the plain Java version is actually faster. Therefore, we can conclude that there is no or no significant performance loss when using instance pointcuts, while the program quality increases.

6. Related Work

AO-extensions for improving aspect-object relationships are proposed in several studies. Sakurai et al. [21] proposed Association Aspects. This is an extension to aspect instantiation mechanisms in AspectJ to declaratively associate an aspect to a tuple of objects. In this work the type of object tuples are declared with a `perobjects` clause and the specific objects are selected by pointcuts. This work offers a method for defining relationships between objects. Similar to association aspects, Relationship Aspects [18] also offer a declarative mechanism to define relationships between objects, which are cross-cutting to the OO-implementation. This work focuses on managing relationships between associated objects. Bodden et al. [7] claim that the two above approaches lack generality

and propose a tracematch-based approach. Although the semantics of the approaches are very similar, Bodden et al. combine features of thread safety, memory safety, per-association state and binding of primitive values or values of non-weavable classes. Our approach, also extending AO, differs from these approaches since our aim is not defining new relationships but using the existing structures as a base to group objects together for behaviour extensions. Our approach also offers additional features of composition and refinement.

The “dflow” pointcut [17] is an extension to AspectJ that can be used to reason about the values bound by pointcut expressions. Thereby it can be specified that a pointcut only matches at a join point when the origin of the specified value from the context of this join point did or did not appear in the context of another, previous join point (also specified in terms of a pointcut expression). This construct is limited to restricting the applicability of pointcut expressions rather than reifying all objects that match certain criteria, as our approach does.

Another related field is Object Query Languages (OQL), which are used to query objects in an object-oriented program (e.g., [8]). However OQLs do not support event based querying, which selects objects based on the events they participate in, as presented in our approach. It is interesting to combine instance pointcuts with OQL. For example instance pointcuts can be used as a predicate in OQL expressions, in order to select from usage-pattern-specific object sets. We will explore possible applications in future work.

Typstates [9] allow users to define a state chart for a type. This specifies which states an object of that type can be in and what causes state transitions. Similar to our approach, state transitions are triggered by runtime events. However unlike instance pointcuts, in typstates an object’s state can only change at method calls where the object is the receiver; with instance pointcuts, objects life-cycle phases can be defined more flexibly by referring to any event where the object is in the dynamic context, e.g., passed as argument or result value. The purpose of the typstates approach is to facilitate more powerful invariant checking at compile-time, whereas we provide a mechanism to actually track object sets at runtime. It would be interesting to investigate possibilities for combining both approaches in the future.

7. Conclusion and Future Work

In our work we have presented instance pointcuts, a specialized pointcut mechanism for reifying sets of objects, which have a similar usage history. Our approach provides a declarative syntax for defining events when a relevant usage pattern starts or ends being applied to an object. Instance pointcuts maintain multisets providing a count for objects, which enter and exit the same usage pattern more than once. Instance pointcut sets can be accessed easily and any changes to these sets can also be monitored with the help of automatically created set monitoring pointcuts. The sets can be declaratively composed, which allows reuse of existing instance pointcuts and consistency among corresponding multisets. Finally, we have presented our modular compilation approach for instance pointcuts based on AspectJ and the ALIA4J Language implementation architecture. ALIA4J provided us with the flexibility AspectJ lacked in instance pointcut composition and type refinement.

The syntax and expressiveness of instance pointcuts partially depend on the underlying AO language; this is evident especially in our usage of the AspectJ pointcut language in the specification of events. Since AspectJ’s join points are “regions in time” rather than events, we had to add the “before” and “after” keywords to our add and remove expressions. Thus, compiling to a different target language with native support for events (e.g., EScala [11] or Composition Filters [2], the point-in-time join-point model [16]) would influence the notation of these expressions.

We think the instance pointcut concept is very flexible and can be useful in various applications. It eliminates boilerplate code to a great extent and provides a readable syntax. We believe that the reuse and composition mechanisms offered by instance pointcuts are beneficial for software evolution since they make it easy to create tailored variations according to new requirements.

One relevant field of application for instance pointcuts are design patterns, which are known to be good examples for aspect-oriented programming [13]. Many design patterns exist for defining the behaviour for groups of objects; thus implementing them with our instance pointcuts seems to provide a natural benefit.

In future work, we would also like to combine our Instance Pointcuts approach with our work on “An Adapter-Aware, Non-Intrusive Dependency Injection Framework for Java” [20]. This framework can inject objects into target locations (e.g., instance fields) and at the same time adapt the objects should they not be type compatible with the target site. Currently, the framework offers limited means to select objects for the injection and injection targets; we intend exploring the use of Instance Pointcuts for these purposes.

References

- [1] I. Aracic, V. Gasiunas, M. Mezini, and K. Ostermann. An overview of caesarj. In A. Rashid and M. Aksit, editors, *Transactions on Aspect-Oriented Software Development I*, pages 135–173. Springer-Verlag, Berlin, Heidelberg, 2006.
- [2] L. M. Bergmans and M. Aksit. How to deal with encapsulation in aspect-orientation. In *Proceedings of OOPSLA 2001 Workshop on Advanced Separation of Concerns in Object-Oriented Systems*, 2001.
- [3] C. Bockisch, M. Haupt, M. Mezini, and K. Ostermann. Virtual machine support for dynamic join points. In *Proceedings of the 3rd international conference on Aspect-oriented software development*, AOSD '04, pages 83–92, New York, NY, USA, 2004. ACM.
- [4] C. Bockisch and M. Mezini. A flexible architecture for pointcut-advice language implementations. In *VMIL '07 : proceedings of the 1st Workshop on Virtual Machines and Intermediate Languages for Emerging Modularization Mechanisms*, New York, March 2007. ACM.
- [5] C. Bockisch and M. Mezini. A flexible architecture for pointcut-advice language implementations. In *Proceedings of the 1st Workshop on Virtual Machines and Intermediate Languages for Emerging Modularization Mechanisms*, VMIL. ACM, 2007.
- [6] C. Bockisch, A. Sewe, H. Yin, M. Mezini, and M. Aksit. An in-depth look at ALIA4J. *Journal of Object Technology*, 11(1):1–28, 2012.
- [7] E. Bodden, R. Shaikh, and L. Hendren. Relational aspects as trace-matches. In *Proceedings of the 7th international conference on Aspect-oriented software development*, pages 84–95. ACM, 2008.
- [8] S. Cluet. Designing oql: Allowing objects to be queried. *Information systems*, 23(5):279–305, 1998.
- [9] R. DeLine and M. Fähndrich. Tpestates for objects. In M. Odersky, editor, *ECOOP 2004 - Object-Oriented Programming*, volume 3086 of *Lecture Notes in Computer Science*, pages 465–490. Springer Berlin Heidelberg, 2004.
- [10] E. Gamma, R. Helm, R. Johnson, and J. Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, Reading, Massachusetts, 1995.
- [11] V. Gasiunas, L. Satabin, M. Mezini, A. Núñez, and J. Noyé. Scala: modular event-driven object interactions in scala. In *Proceedings of the tenth international conference on Aspect-oriented software development*, AOSD '11, pages 227–240, New York, NY, USA, March 2011. ACM.
- [12] Google. Caliper: Microbenchmarking framework for java. <http://code.google.com/p/caliper/>, 2014.
- [13] J. Hannemann and G. Kiczales. Design pattern implementation in Java and AspectJ. In *Proceedings of the 17th ACM conference on Object-oriented programming, systems, languages, and applications*, pages 161–173. ACM Press, 2002.
- [14] F. Heidenreich, J. Johannes, M. Seifert, and C. Wende. Closing the gap between modelling and java. In M. van den Brand, D. Gašević, and J. Gray, editors, *Software Language Engineering*, volume 5969 of *Lecture Notes in Computer Science*, pages 374–383. Springer Berlin / Heidelberg, 2010. 10.1007/978-3-642-12107-4_25.
- [15] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. Griswold. An overview of aspectj. *ECOOP 2001 Object-Oriented Programming*, pages 327–354, 2001.
- [16] H. Masuhara, Y. Endoh, and A. Yonezawa. A fine-grained join point model for more reusable aspects. In N. Kobayashi, editor, *Programming Languages and Systems*, volume 4279 of *Lecture Notes in Computer Science*, pages 131–147. Springer Berlin / Heidelberg, 2006. 10.1007/11924661_8.
- [17] H. Masuhara and K. Kawauchi. Dataflow pointcut in aspect-oriented programming. In *In APLAS03 - the First Asian Symposium on Programming Languages and Systems*, pages 105–121, 2003.
- [18] D. Pearce and J. Noble. Relationship aspects. In *Proceedings of the 5th international conference on Aspect-oriented software development*, pages 75–86. ACM, 2006.
- [19] H. Rajan and K. J. Sullivan. Eos: instance-level aspects for integrated system design. In *Proceedings of the 11th ACM SIGSOFT Symposium on Foundations of Software Engineering 2003 held jointly with 9th European Software Engineering Conference, ESEC/FSE 2003*, pages 291–306, 2003.
- [20] A. Roemers, K. Hatun, and C. Bockisch. An adapter-aware, non-intrusive dependency injection framework for java. In *Proceedings of the International Conference on Principles and Practices of Programming on the Java Platform: virtual machines, languages and tools*, PPPJ '13. ACM, 2013.
- [21] K. Sakurai, H. Masuhara, N. Ubayashi, S. Matuura, and S. Komiya. Design and implementation of an aspect instantiation mechanism. *Transactions on aspect-oriented software development I*, pages 259–292, 2006.
- [22] E. Tanter. Expressive scoping of dynamically-deployed aspects. In *Proceedings of the 7th international conference on Aspect-oriented software development*, AOSD '08, pages 168–179, New York, NY, USA, 2008. ACM.