Language Features for Software Evolution and Aspect-oriented Interfaces: An Exploratory Study[*]

Robert Dyer[1], Hridesh Rajan[1], and Yuanfang Cai[2]

[1] Iowa State University
Dept. of Computer Science
`{rdyer,hridesh}@iastate.edu`
[2] Drexel University
Dept. of Computer Science
`yfcai@cs.drexel.edu`

Abstract. A variety of language features to modularize crosscutting concerns have recently been discussed, e.g. open modules, annotation-based pointcuts, explicit join points, and quantified-typed events. All of these ideas are essentially a form of aspect-oriented interface between object-oriented and crosscutting modules, but the representation of this interface differs. Previous works have studied maintenance benefits of AO programs compared to OO programs, by usually looking at a single AO interface. Other works have looked at several AO interfaces, but only on relatively small systems or systems with only one type of aspectual behavior. Thus there is a need for a study that examines large, realistic systems for several AO interfaces to determine what problems arise and in which interface(s). The main contribution of this work is a rigorous empirical study that evaluates the effectiveness of these proposals for 4 different AO interfaces by applying them to 35 different releases of a software product line called Mobile-Media and 50 different releases of a web application called Health Watcher. In total over 400k lines of code were studied across all releases. Our comparative analysis using quantitative metrics proposed by Chidamber and Kemerer shows the strengths and weaknesses of these AO interface proposals. Our change impact analysis shows the design stability provided by each of these recent proposals for AO interfaces.

1 Introduction

It is generally accepted that advanced separation of concerns tehcniques give software engineers new and valuable possibilities to organize concerns in their system and improve its overall modularity [20, 21, 23, 25, 30]. Early work in this area was on programming language features that focussed on separating base concerns and crosscutting concerns. However, since early 2000 some consensus has emerged in the research community that a notion of interface between base concerns and now separated, but

previously crosscutting, concerns may be necessary to truly realize the full potential of advanced separation of concern techniques [1, 18, 23, 30, 36, 39].

A number of researchers have responded to this observation. For example, Aldrich proposed the notion of Open Modules [1] that make the base code aware of the join points possibly advised, by writing module definitions that state what join points are available for advising by aspects. Thus if a join point changes, the maintainer must also update the module definition which alerts the aspect maintainer to verify their pointcuts.

Around the same time, Kiczales and Mezini proposed the notion of explicit join point matching using annotations [23]. In their proposal, join points are explicitly marked using Java annotations and the pointcuts target the annotation instead of using an implicit pattern, avoiding the fragility of implicitly matching based on names.

More recently, Steimann *et al.* [36], Hoffman and Eugster [18], and Rajan and Leavens [30] have explored providing a notion of explicit join points with quantification using types. The basic difference between these ideas and the notion of matching using annotations is two fold. First, arbitrary program points can be marked as a join point, whereas annotation-based marking is limited to interface level join points. Second, pointcuts match using a declared type that gives information about the context available at the join point and can also be used for selecting all join points of that type. This type acts as an interface between base and aspect code [18, 30, 36].

All of these proposals give clear, compelling and representative examples to demonstrate their main ideas. These examples typically focus on the problem at hand such that it is easy to demonstrate the issue and the benefits of the proposed solution. What is not clear is: to what extent do the claimed benefits of each of these ideas help developers, by decreasing the impact of change in a software system?

1.1 Case Study Overview

There has been a large body of recent case studies on the software engineering (SE) benefits of aspect-orientation [12, 13, 18, 23, 24]. These works compute standard SE metrics such as coupling and cohesion and compare aspect-oriented (AO) designs to object-oriented (OO) designs or use the metrics to determine stability and fault-proneness of the systems. However, most of these works focus on comparing AspectJ [20] to Java and do not compare different AO interfaces with each other, leaving developers to wonder about the benefits of one proposal over others.

This work helps fill that gap by studying and comparing different proposals for aspect-oriented interfaces to investigate how these interfaces impact code changes. For this, we consider a software product line for handling multimedia on mobile devices, called MobileMedia [13] and a web-based health application, called Health Watcher [24, 34]. MobileMedia is a medium sized system with 8 releases. Health Watcher is also a medium sized system with 10 releases. In total, this study contains over 400k lines of code across these 18 releases.

Similar to previous in-depth analyses by Figueiredo *et al.* [13] and Greenwood *et al.* [16], we present metrics such as coupling and cohesion as well as an analysis of the change propagation across releases. However, unlike those studies we consider not only OO and pattern-based pointcuts (PCD) but also three other proposals for AO in-

terfaces: open modules [1] (OM), annotation-based pointcuts [23] (@PCD), and quantified, typed events [30] (EVT).

1.2 Results and Contributions

There were several interesting results to come out of our case study. First, the annotation-based pointcut and quantified, typed event approaches showed several benefits, in terms of change impact, over the standard pattern-based pointcut approach.

- The @PCD releases have 18% fewer changed pointcuts than the PCD releases, due to a lack of fragile pointcuts.
- The total number of changed event types in MobileMedia is 74% fewer than the total number of changed pointcuts in the PCD releases and 66% fewer than the total pointcuts changed in the @PCD releases.

Second, the PCD and @PCD releases showed benefit over EVT for certain design rules.

- For the EVT releases, we had to be aware of and manually maintain design rules related to encapsulating entire types (e.g. to make an entire class synchronized). The PCD, @PCD, and OM releases used pointcuts to automatically maintain such design rules.
- Such design rules show cases where patterns do not exhibit fragile pointcut behavior, as the pointcuts are expected to capture all methods in the advised types.

Additionally, the EVT releases showed some benefit over the @PCD releases due to its ability to uniformly access context information when announcing events.

In summary, the key contributions of the case study performed in this work are:

- The first rigorous study of different language features for four different AO interfaces on substantial case studies.
- A suite of tools to automate measuring change propagation for PCD, OM, @PCD, and EVT. This automation reduces the chance for errors in our empirical study. These tools are released in the public domain, thus they will be useful to empirical researchers conducting studies of a similar nature.
- A new set of 21 MobileMedia and 30 Health Watcher releases using @PCD, OM, and EVT interfaces. These artifacts are also released in the public domain with the hope that they will encourage additional rigorous measurements in aspect-oriented language research.
- A change propagation analysis, that shows the stability gained from designs using annotation-based pointcuts and quantified, typed events in the face of fragile pointcuts [13, 30, 37].

Next we describe some prior studies on AO interfaces. In Section 3 we introduce the studied language designs. We then present our case study in Sections 4–7. Then we conclude with discussion in Section 9 and future work.

2 Related Work

Previous case studies on AO interfaces can be broken down into two categories: comparative studies that compare one or more AO language feature with standard OO features and studies that focus solely on measuring metrics to predict maintainability of an AO language feature. Here, we describe some of these studies.

2.1 Language Feature Comparison Studies

Figueiredo *et al.* [13] studied the effects of evolving software product lines (SPLs) using aspects. Similar to our study, they measure change propagation and a set of standard metrics (such as coupling and cohesion). Their study showed some of the pros and cons to using AO language features when compared to OO features. For example, their study showed that changes affecting core features (such as changing a mandatory feature into an optional feature) are not well suited for AO. However, their study was limited to only one AO interface (pattern-based pointcuts).

Hoffman and Eugster [18] studied the coupling, cohesion and separation of concerns for several projects with implementations in Java, AspectJ, and explicit join points (EJPs). Their study focused solely on implementing exception handling with each AO interface. Similar to our study, their study examines software engineering metrics and compares each AO interface against each other. Our study however looks at a total of 4 AO interfaces and multiple types of crosscutting behavior (instead of just exception handling) in two distinct systems with a total of 68 AO releases.

Kiczales and Mezini [23] studied seven different AO interfaces for improving separation of concerns in AspectJ-like languages. These included standard method calls, explicit join points using annotation-based pointcuts and implicit pattern-based pointcuts. They analyze each mechanism based on locality, explicit/implicit and ease of evolution and then provide guidelines on when each mechanism should be used in practice. Our work is similar in the sense that we analyze several language interfaces. Their work uses a simple example for comparison while our work examines the medium-sized MobileMedia [13] software product line and Health Watcher [16, 34] web application.

2.2 Maintenance Studies

Ferrari *et al.* [12] studied several SPLs to determine the possible language features that led to faults in those systems. Their results show that obliviousness was a key cause of faults in those systems and that pattern-based pointcuts are not necessarily the main cause of faults in AO designs. Their study determined the cause of faults in AO systems while our study examines the effects of several AO interfaces on software maintenance.

Kulesza *et al.* [24] investigated the effect of AO interfaces on software maintenance by measuring standard software engineering metrics. They measured separation of concerns, coupling, cohesion, and size and concluded that in the presence of widely-scoped design changes, the AO designs exhibited superior stability and reusability compared to OO designs. In their study, they look at 2 releases of the Health Watcher application. Our study on the other hand examines 10 releases of Health Watcher and 7 releases of MobileMedia, giving us more variability to examine and allowing us to analyze the

effects of varying types of interfaces added to a system. Their work also focuses solely on pattern-based pointcuts, whereas we consider several AO interfaces.

Similar to Hoffman and Eugster, Filho *et al.* [14] studied how implementations of exception handling in Java and AspectJ compared. Again, their study mostly focused on one type of crosscutting behavior and only examined a single AO interface.

2.3 Summary

Prior works tend to focus on a single AO interface in their comparison (usually pattern-based pointcuts) and compare AO designs to a similar OO design. The studies that do compare more than one AO interface (such as Kiczales and Mezini [23]) tend to use small examples (4 classes) or only look at one type of modularization (such as Hoffman and Eugster [18]). Thus there is a need for a case study that examines both larger systems and multiple AO interfaces.

3 Background

In this section, we describe our language feature selection process and then give an overview of each studied AO interface using an example based on a pattern occurring frequently in one of our case study candidates, MobileMedia [13].

3.1 Language Feature Selection

Quantification is how aspect-oriented techniques select points in the program to provide additional behavior. Different languages provide different features for controlling quantification. For this study, we categorized the various AO interfaces in recent literature and determined three main categories based on how quantification is controlled.

- In the first category, quantification is controlled solely by aspects.
- In the second category, quantification is controlled solely by the base code.
- In the third category, quantification is controlled by an intermediary between base components and aspects.

Language	Controls Quantification	Implementation?	Selected?
Pattern-based pointcuts [21]	Aspects	Yes	**Yes**
Aspect-aware Interfaces [22]	Base Code	No	No
Open modules [1, 26]	Base Code	Yes	**Yes**
Annotation-based pointcuts [23]	Intermediary	Yes	**Yes**
IIIA [36]	Intermediary	Yes	No
Join point interfaces [19]	Intermediary	Yes	No
Quantified, typed events [30]	Intermediary	Yes	**Yes**
XPIs [39]	Intermediary	No	No

Fig. 1. Languages, their features, and our selection (this list is not exhaustive)

For the first category, pattern-based pointcuts [21] (PCD) was the most relevant choice, as it is used in industry and also highly researched.

For the second category, there were two possible candidates: aspect-aware interfaces [22] and open modules [1]. We selected open modules (OM) because an implementation [26] was readily available.

For the third category, there were several candidates: XPIs [39], annotation-based pointcuts [23], implicit invocation with implicit announcements [36] (IIIA), join point interfaces [19] (JPIs), and quantified, typed events [30].

We selected quantified, typed events (EVT) due to our familiarity with its compiler infrastructure. Since one of the authors was involved in the design of this language, we decided to pick a second candidate in this category and chose annotation-based pointcuts (@PCD), as this language feature was not developed by the authors and once again compiler support was readily available.

It is important to note that we only selected a handful of candidates and while we grouped them into categories for our selection process, the results of our study may not generalize to other candidates even within the same category as each language has varying features. Still, this categorizing was a useful mechanism for us to select a subset of varying language features for our study.

3.2 Running Example: Exception Handling

Throughout this section, we use the same example which was taken directly from the MobileMedia case study candidate. Note that this example may not represent the best possible design for any particular language. Instead, it represents a real-world example and helps to highlight some of the potential issues that can arise when using the various AO interfaces.

Let us consider the class `FileScreen` shown in Figure 2. This class represents a screen presented to a user for manipulating a file. When the `saveCommand` (line 5) is requested, the class saves the data to the specified file name. When the `deleteCommand` (line 11) is requested, the file is deleted. The screen is shown on a `display` (line 2), which can be updated to show different screens (line 9).

An example requirement for such a class is to consistently display error messages to the user. There may be multiple screens that deal with I/O and all such screens should consistently handle errors that occur during that I/O by showing the I/O error screen. Note that in some cases, the designers have decided no error should be displayed (for example, when deleting a file and it was deleted by another user between the time of request and handling of the command).

Notice that this requirement would thus be scattered across multiple screen classes. Thus this is an example of a crosscutting concern and aspect-oriented [21] interfaces can help modularize this concern. Next we examine how four such AO interfaces can achieve such a modularization.

3.3 Exception Handling Using Pattern-Based Pointcuts

The aspect `ExceptionHandler` shown in Figure 3 (lines 12–27) implements the requirement to consistently handle all exceptions using pattern-based pointcuts. This

```
1  class FileScreen {
2    Display display;

4    void handleCommand (Command c) {
5      if (c == saveCommand) {
6        try {
7          // open the file and save data
8        } catch (FileNotFoundException e) {
9          display.ShowFileError (e);
10       }
11     } else if (c == deleteCommand) {
12       try {
13         // delete the file
14       } catch (FileNotFoundException e) {
15         // do nothing
16       }
17     }
18   }
19 }
```

Fig. 2. Exception Handling with OO features (based on a code pattern from MobileMedia [13])

aspect contains an around advice (lines 15–20), which when triggered will properly handle the exception. The named pointcut savepc (lines 13–14) matches the execution of the method save, which had to be created in order to have a join point capable of being advised by the aspect. This is an example of quantification failure [39]. More details are given later in Section 7.4. The same problem exists for the delete method.

The pointcut is also fragile [13, 30, 37], due to using the name of the save method. If the method is renamed inside the class, the aspect will no longer match that point and the pointcut must also be updated to reflect this renaming. More details are given later in Section 7.5. Again, this problem exists for the deletepc pointcut as well.

The advice uses the display variable, which is not exposed as context in the pointcut and is instead accessed indirectly through available context (the receiver object, screen). This inability to express the exact context information needed introduces unintended coupling to the receiver object's class. More details are given later in Section 7.6.

3.4   Exception Handling Using Quantified, Typed Events

Quantified, typed events [30] allow programmers to declare named event types. An *event type declaration* p has a return type, a name, and zero or more context variable declarations. These *context declarations* specify the types and names of reflective information communicated between announcement of events of type p and handler methods registered for announcement of events of type p. These declarations are independent from the modules that announce or handle these events. The event types thus provide an interface that completely decouples subjects and observers. An example event type declaration is shown in Figure 4 (line 1). The **event** FileSaveEvent declares that events of this type make one piece of context available: the display.

```
1 class FileScreen {
2   Display display;

4   void handleCommand (Command c) {
5     if (c == saveCommand) { save (); }
6     else if (c == deleteCommand) { delete (); }
7   }
8   void save () { /* open the file and save data */ }
9   void delete () { /* delete the file */ }
10 }

12 aspect ExceptionHandler {
13   pointcut savepc(FileScreen screen):
14     execution(* FileScreen.save ()) && this(screen) {
15   around(FileScreen screen): savepc(screen) {
16     try { proceed (); }
17     catch (FileNotFoundException e) {
18       screen.display.ShowFileError (e);
19     }
20   }
21   pointcut deletepc(FileScreen screen):
22     execution(* FileScreen.delete ()) && this(screen) {
23   around(FileScreen screen): deletepc(screen) {
24     try { proceed (); }
25     catch (FileNotFoundException e) { /* do nothing */ }
26   }
27 }
```

Fig. 3. An example usage of pattern-based pointcuts [20]

The class `FileScreen` (lines 4–14) declares and announces an event of type `FileSaveEvent` using an *announce statement* [30] (line 9). Arbitrary blocks can be declared as the body of an announce statement, which avoids quantification failure. These blocks can be replaced by handler methods, giving functionality similar to around advice in pattern-based languages. The event type `FileSaveEvent` declares one context variable, thus the announce statement binds the field `display` to the context variable named `display` (line 9).

Finally, the names of event declarations can be utilized for quantification in a binding declaration. A *binding declaration* [30], binding in short, associates a handler method to a set of events identified by an event type. The binding (line 25) says to run the method **handler** when events of type `FileSaveEvent` are announced. This allows quantifying over all announcements of `FileSaveEvent` with a succinct binding declaration, without depending on the modules that announce those events. Use of event names in bindings simplifies them and avoids coupling the observers to the subjects.

Despite the name, quantified, typed events follow a unified model [31, 32] and do not actually distinguish between classes containing handlers and normal classes. For ease of reading however, we call any class containing a handler method a handler class.

Handler classes also contain *register statements* [30] (line 17). These statements make all bindings contained within the handler class active for the instance and allow for dynamically adding or removing advice functionality in the system. Pattern-based pointcuts can emulate this functionality by adding a boolean flag to aspects.

```
1  void event FileSaveEvent { Display display; }
2  void event FileDeleteEvent { Display display; }

4  class FileScreen {
5    Display display;

7    void handleCommand (Command c) {
8      if (c == saveCommand) {
9        announce FileSaveEvent(display) { /* open the file and save data */ }
10     } else if (c == deleteCommand) {
11       announce FileDeleteEvent(display) { /* delete the file */ }
12     }
13   }
14 }

16 class ExceptionHandler {
17   ExceptionHandler() { register (this); }

19   void handler(FileSaveEvent next) throws Throwable {
20     try { next.invoke(); }
21     catch (FileNotFoundException e) {
22       next.display().ShowFileError (e);
23     }
24   }
25   when FileSaveEvent do handler;

27   void handler(FileDeleteEvent next) throws Throwable {
28     try { next.invoke(); }
29     catch (FileNotFoundException e) { /* do nothing */ }
30   }
31   when FileDeleteEvent do handler;
32 }
```

Fig. 4. An example usage of quantified, typed events [30]

Each handler method takes an event closure as the first argument. An *event closure* [30] contains code needed to run other applicable handlers and the original event's code. An event closure is run by an *invoke expression*. The **invoke** expression in the implementation of the handler method (line 20) causes other applicable handlers and the original event's code to run before handling any exceptions.

While this version of the program did not require refactoring to expose join points, it still requires modifying the class in order to add explicit event announcements. Unlike the pattern-based version which only modifies the class in cases where it needs to refactor to expose a join point, the Ptolemy version must modify each class to add explicit announcements.

3.5 Exception Handling Using Annotation-based Pointcuts

When using pattern-based pointcuts, the code being advised by aspects is completely oblivious to those aspects. One approach that is less oblivious, which is quite similar looking to quantified, typed events, is to mark each advised join point with an annotation [23]. The aspects then match based on that annotation.

```
1  @interface FileSaveEvent { }
2  @interface FileDeleteEvent { }

4  class FileScreen {
5    Display display;

7    void handleCommand (Command c) {
8      if (c == saveCommand) { save (); }
9      else if (c == deleteCommand) { delete (); }
10   }
11   @FileSaveEvent
12   void save () { /* open the file and save data */ }
13   @FileDeleteEvent
14   void delete () { /* delete the file */ }
15 }

17 aspect ExceptionHandler {
18   pointcut savepc(FileScreen screen):
19       execution(@FileSaveEvent * *(..)) && this(screen) {
20   around(FileScreen screen): savepc(screen) {
21     try { proceed (); }
22     catch (FileNotFoundException e) {
23       screen.display.ShowFileError (e);
24     }
25   }
26   pointcut deletepc(FileScreen screen):
27       execution(@FileDeleteEvent * *(..)) && this(screen) {
28   around(FileScreen screen): deletepc(screen) {
29     try { proceed (); }
30     catch (FileNotFoundException e) { /* do nothing */ }
31   }
32 }
```

Fig. 5. Exception Handling with annotation-based pointcuts [23] (changes from Figure 3 are highlighted)

For example, consider the code shown in Figure 5. The gray portions of the code represent what was changed from the pattern-based implementation. An annotation FileSaveEvent was created (line 1) and then used to mark the advised join point method save (line 12). The pointcut for the aspect (line 14) was modified to become **execution**(@FileSaveEvent * *(..)) and match based on that annotation instead of matching against the string representation of the method name.

Similar to quantified, typed events, this AO interface helps avoid the fragile pointcut problem: if the method save is renamed the pointcut will still match (to the annotation). This comes as a trade-off of limiting the obliviousness of the base code. One may argue that loss of obliviousness is not necessarily a trade-off, as it is not clear this property is actually desirable [33, 35]. Regardless, explicitly marking join points does not solve every problem.

For example, the quantification failure problem still exists when explicitly marking join points. We still had to refactor the code to create the save method in order to annotate and advise the code.

This release also suffers from problems passing context to the advice. The `display` field is not available using the standard pointcuts (this, target, args). Instead, we were forced to access the field through the receiver object (`screen`), which makes the aspect coupled to the interface of both `FileScreen` and `Display`. The quantified, typed event release does not suffer from this problem.

3.6 Exception Handling Using Open Modules

Open modules [1] declare which join points are exposed to aspects via a module definition. This puts the burden onto the module maintainer to maintain relationships between join points in the base code and pointcuts matching those points.

Ongkingco *et al.* [26] proposed an extended version of open modules and an implementation for AspectJ [20]. We use their implementation's syntax for this example. Figure 6 is an example module for our exception handling requirement. The figure omits the class `FileScreen` and aspect `ExceptionHandler`, as they are identical to Figure 3.

```
1 /* class FileScreen and aspect ExceptionHandler same as in Figure 3. */
2 module Exceptions {
3   class FileScreen;
4   expose : ExceptionHandler.savepc(FileScreen);
5   expose : ExceptionHandler.deletepc(FileScreen);
6 }
```

Fig. 6. Exception handling with open modules [1, 26]

The module `Exceptions` (lines 2–6) for the class `FileScreen` (line 3) exposes two named join points in that class (lines 4–5): the pointcuts `savepc` and `deletepc` defined in the aspect `ExceptionHandler`. The module states that the aspect is allowed to match these join points.

One maintenance issue occurs if the signature of the pointcuts change in the aspect, as the maintainer of the module(s) would be required to update the module definition(s) as well. For example, if the context information passed changes types then both the pointcut in the aspect and the pointcut in the module would need to update accordingly.

Open modules also suffer from limited availability of context information. Once again the context made available is the `FileScreen` and not the actual context needed (the display). This is similar to the pattern-based and annotation-based approaches.

4 Case Study Overview

To evaluate the proposed AO interfaces studied here, we examined them in the context of two applications: an existing software product-line application called MobileMedia [13] and an existing web application called Health Watcher [16, 34]. This section

describes our experimental setup, the technique used to generate new releases of the studied applications and the tools developed and used for the study.

In order to perform this case study, we created a total of 51 modified releases of the MobileMedia and Health Watcher applications, modified 2 compilers to automatically compute various software engineering metrics and created a tool for automatically measuring change propagation. All artifacts and tools are available for download[3]. An important advantage of these tools was that they removed the manual, and often error-prone, steps from our empirical study.

4.1 New Code Artifacts

The OO and pattern-based pointcut code artifacts for this study were re-used from previous work [13, 16, 34]. We created the artifacts using annotation-based pointcuts, open modules, and quantified, typed events since they did not previously exist for either application. When creating these artifacts, *our objective was to keep other variables such as design strategy constant between all versions and only change the crosscutting feature*. Figure 7 gives an overview of the artifacts re-used and created for our study.

	OO	*PCD*	*@PCD*	*OM*	*EVT*
MobileMedia	Existing [13]	Existing [13]	**New**[1]	**New**[2]	**New**
Health Watcher	Existing [16, 34]	Existing [16, 34]	**New**[1]	**New**[2]	**New**

Fig. 7. Overview of all code artifacts studied. Note 1: Based on recommendations of Kiczales and Mezini [23]. Note 2: Based on recommendations of Ongkingco *et al.* [26].

For example, starting with release 4 of the Health Watcher releases for AspectJ, an observer pattern aspect library was used. This library was re-used in the annotation-based pointcut, open modules and Ptolemy releases despite the fact that Ptolemy's quantified, typed events actually make this library unnecessary (the events implement the observer pattern directly, so no library is needed). Removing this library however would change the base and aspect components in the system and introduce extra variables into the analysis. Leaving it in place meant the only difference between the Ptolemy and other releases was the quantification mechanism used for implementing crosscutting behavior.

Creating Annotation-based Pointcut Releases Using the pattern-based pointcut releases as a starting point, we implemented all 7 releases of MobileMedia and all 10 releases of Health Watcher using an annotation-based pointcut syntax [23]. We modified each pointcut in every aspect to match based on a new annotation and for each join point in the base code matching the original pointcut, we annotated the method with the new annotation.

---

[3] **Tools/artifacts download:**
http://ptolemy.cs.iastate.edu/design-study/

For example, consider Figure 8 which shows a pattern-based aspect on the left. For each pointcut pattern, we generate an interface (lines 1 & 2). We then modify the pointcut to match the annotation (lines 5 & 9) and then annotate the relevant types or methods so the set of matched points remains the same.

```
                                    1 @interface E1 { }
                                    2 @interface E2 { }

1 aspect A {                        4 aspect A {
2   after() : execution(pcd1) {     5   after() : execution(@E1 * *.*(..)) {
3     // advice1                     6     // advice1
4   }                               7   }

6   around() : execution(pcd2) {    9   around() : execution(@E2 * *.*(..)) {
7     // advice2                     10    // advice2
8   }                               11  }
9 }                                 12 }
```

Fig. 8. Creating annotated pointcut releases. Points in base code matching the pointcuts `pcd1` and `pcd2` are annotated with the new annotations `E1` and `E2`.

The names of the annotations were chosen based on properties of the code, following the guidelines of Kiczales and Mezini [23]. They state that *"[w]hen using named attributes, choose a name that describes what is true about the points, rather than describing what a particular advice will do at those points."* [23, p.207] The results were verified by comparing the weaving logs produced by the standard AspectJ compiler (ajc) for both the original pattern-based pointcut releases and the new annotation-based pointcut releases.

```
1 aspect A {                       1 module M {
2   after() : pkg.C.pcd1() {       2   class pkg..*;
3     // advice1                    3   expose : pkg.C.pcd1();
4   }                              4   expose : pkg.D.pcd2();
                                   5 }
6   around() : pkg.D.pcd2() {
7     // advice2
8   }
9 }
```

Fig. 9. Creating open modules releases

Creating Open Module Releases To study the effect of open modules [1], we implemented all 7 MobileMedia and all 10 Health Watcher releases using the AspectJ-based implementation of open modules [26]. Starting with the first release, we made a copy of the pattern-based pointcut release and then created module definitions. Ongkingco *et al.* state *"the module hierarchy is envisioned to closely follow the package*

hierarchy." [26, p.6] We follow this recommendation and created modules to follow the package structure of the system.

For example, consider the aspect shown on the left side of Figure 9. This aspect uses pointcuts defined in classes C and D. A module definition is created that includes all classes in the package pkg. This module exposes the pointcuts defined in the classes C and D, thus making them available for the aspect to advise.

Since the only difference between the pattern-based and the open modules releases are the addition of module definitions, for each subsequent release we started by copying the pattern-based pointcut release. Then we copied the module definition(s) from the previous open modules release. Modules were then updated to reflect changes in the base code and, where appropriate, new modules were added.

Creating Quantified, Typed Event Releases For each quantified, typed event [30] release we started with the pattern-based pointcut release as a template, creating one handler class for each aspect. For each advice body in an aspect, a new handler method was added to the handler class. Event types were created and event announcement added to emulate the pattern-based pointcut-advice semantics.

```
1 void event E1 { .. }
2 void event E2 { .. }

1 aspect A {                          4 class A {
                                      5   A() { register (this); }

2   after() : pcd1() {                7   void handler(E1 next) throws Throwable {
3     // advice1                       8     // advice1
4   }                                  9   }
                                      10   when E1 do handler;

5   around() : pcd2() {               12   void handler(E2 next) throws Throwable {
6     // advice2                      13     // advice2
7   }                                 14   }
                                      15   when E2 do handler;
8 }                                   16 }
```

Fig. 10. Creating quantified, typed event releases. Points in base code matching the pointcuts `pcd1` and `pcd2` are annotated with event announcements for events `E1` and `E2`.

For example, consider the aspect shown in the left side of Figure 10, which contains 2 pieces of advice. This aspect is translated into the class on the right side. For each piece of advice, an event type declaration is generated (lines 1 & 2). Then at each point selected by the original pointcut pattern, an event announcement was added for the event type. Each advice body is translated into a handler method (lines 7 & 12) with a matching binding declaration (lines 10 & 15).

Note that since the initial work on Ptolemy [30], the language has been extended to include support for inter-type declarations. The syntax is identical to that of AspectJ and the implementation was directly borrowed from the ABC AspectJ compiler [2] and

added to the Ptolemy compiler, as the research version of the Ptolemy compiler is also based on the JastAdd [10] extensible compiler framework.

4.2 Automation of Empirical Evaluation

Evaluating the benefits of the studied designs using standard software engineering metrics and change propagation by hand can be tedious and error-prone. To solve this problem, we built several tools to automatically measure these metrics and allow for consistency. These included several modified compilers and tools for measuring change propagation. Our tool support builds on the open-source ABC [2] AspectJ compiler. The ABC compiler was used for two reasons: it has a JastAdd [10] extensible frontend available which simplifies extensions and it contains support for the only known implementation of open modules. The use of ABC was also driven by the fact that the research version of the Ptolemy compiler is also JastAdd-based and our tool extensions could be re-used for both compilers (giving us automated tool support for every studied language design).

Measuring Change Propagation To measure change propagation, a JastAdd module was created to serialize the parsed AST into an XML format. Since every compiler used in our study is based on the JastAdd extensible compiler, the new functionality was shared as a reusable module between these compilers. This also ensured that change propagation measurement was done consistently.

A separate tool was created that takes two of these XML files as input, representing two versions of the same code tree, and compares the two trees to determine which components are new, were removed, or have changed. We considered a renamed component (including moving it to another package) as a change (instead of a remove and an add) and manually identified such renames in a separate XML file to aid the tool.

The tool is capable of determining changes at the granularity of classes, aspects, event types, annotations, and pointcuts. The results were then manually verified against diffs of the MobileMedia code releases for Java and AspectJ to ensure the accuracy of the tool.

Measuring Software Engineering Metrics To measure the metrics suite proposed by Chidamber and Kemerer [5] for coupling and cohesion, we created another JastAdd module which measures and reports these metrics. This module was shared and used in each compiler in our study.

Chidamber and Kemerer propose that a component is coupled to another component if it accesses a field or calls a method from the other component. They also propose a class is cohesive if the operations of the class operate on similar attributes of the class.

We used the previous results from Figueiredo *et al.* [13] as a guide for our implementation, comparing the values for the OO and pattern-based pointcut MobileMedia releases to their previously published results. The formalization of these metrics are given in the previous works [5, 13].

No extension to the metrics was necessary for the annotation-based pointcuts, as the existing OO and pattern-based pointcut metrics apply directly. Note that while it

is not obvious, the compiler must check the interface of annotations (specifically, the `@Target` annotation) in order to verify they can be applied to the specified element. Thus, this is treated like a field access and any mention of an annotation adds coupling to that annotation's type.

The metrics suite was extended to support open modules and quantified, typed events in a straight-forward manner, similar to the pattern-based pointcut extensions [13]. Five different metrics used in our study were extended: number of components (NOC), number of attributes (NOA), number of operations (NOO), coupling between components (CBC), and lack of cohesion in operations (LCOO).

For open modules releases, modules are treated like classes and join point exposures treated similar to an aspect pointcut. Thus, the metrics are extended as follows:

– NOC: each module adds 1 to this metric
– NOA: for each module, for each join point exposure, 1 is added to this metric
– NOO, CBC[4], LCOO: no extension necessary

For quantified, typed events, event type declarations are treated like a class with context considered a field of the event type. Announcing an event is treated like a method call. Thus, the metrics are extended as follows:

– NOC: each event type adds 1 to this metric
– NOA: for each event, for each declared context, 1 is added to this metric
– NOO: since handler methods are normal methods, 1 is added to this metric
– CBC: any handler class that mentions an event type and any base class that announces an event type is coupled to that type and adds 1 to this metric
– LCOO: the metric is computed the same, however since handler methods are normal methods in a class they are also included in this metric's computation

4.3 Threats to Validity

In this section we discuss internal and external threats to the validity of our case study.

Internal Validity To reduce the risk of bias when selecting languages for study, we first decided the focus of the study to be examining the effect of AO interfaces for minimizing pointcut fragility and change propagation. Then we categorized existing AO interfaces by how they achieve quantification (see Section 3.1). As one of the selected languages (Ptolemy) was designed by some of the authors, we also selected a second candidate (annotated pointcuts) from that category.

The code artifacts created for this study were the MobileMedia and Health Watcher (releases for annotation-based pointcuts (@PCD), open modules (OM), and quantified, typed events (EVT). To reduce the risks associated with creating these artifacts, we attempted to keep other variables constant (such as design strategy used) and only vary the quantification mechanism used. See Section 4.1 for more details.

---

[4] Note that CBC only measures *explicit* coupling and thus (similar to pointcuts) the modules's join point exposures do not affect this metric.

We reduced the risk associated with creating the EVT, @PCD, and OM releases by first basing them off the existing pattern-based pointcut releases (which were not created by any of the authors). Next, we used recommendations by experts in each respective language in their published work [23, 26, 30] to modify the pattern-based pointcut releases and create the releases for the new AO interfaces.

For example, we followed the guidelines given by the implementers of the open modules implementation used to create one module definition for each package [26]. We also followed a naming scheme proposed by Kiczales and Mezini [23] when generating annotations for the @PCD releases, which was shown to offer design stability. See Section 4.1 for more details.

External Validity Regarding external validity, we identified a threat that the studied systems may not faithfully represent software in industry. This risk is reduced since the applications are implemented in both Java and AspectJ, which is a representative approach in the AO domain. Further, MobileMedia is a software-product line comprised of 8 releases based on industry-strength technologies for mobile systems, such as the Java Mobile Information Device Profile (MIDP) and Mobile Media API (MMAPI). Additionally, this system has been studied extensively [12–14].

Similarly, Health Watcher is a real-world application used for reporting health complaints. This system uses several industrial strength technologies/techniques, such as persistence mechanisms, remote invocation (RMI), concurrency, JDBC, etc.

Another external threat is regarding the generalizability of our study. The results shown in our study clearly demonstrate the issues with each language feature studied, however these results may not generalize to other language features or to languages that contain more than one of these features.

5 Case Study: MobileMedia

This section contains our first studied project, a software product-line application called MobileMedia [13].

5.1 MobileMedia Overview

MobileMedia is an extension of MobilePhoto [40], which was developed to study the effect of AO designs on software product lines (SPL). MobileMedia is an SPL for applications that manipulate photos, music, and videos on mobile devices. MobileMedia extends MobilePhoto to add new mandatory, optional and alternative features.

There are a total of 8 releases and descriptions of each are shown in Figure 11. For example in release 7 (R7) a new feature is added to manage music and a feature added in a previous release to manage photos is turned into an alternative feature. Note that release 1 is the same across all languages studied, and thus omitted from discussion in this paper.

We chose to study MobileMedia for the following reasons.

1. It was successfully used in several previous AO studies [12–14].

Release	Description	Type of Change
R1	MobilePhoto core	
R2	Exception handling included (in the aspect-oriented and Ptolemy releases, exception handling was implemented according to [14])	Inclusion of non-functional concern
R3	New feature added to count the number of times a photo has been viewed and sorting photos by highest viewing frequency. New feature added to edit the photo's label	Inclusion of optional and mandatory features
R4	New feature added to allow users to specify and view their favorite photos.	Inclusion of optional feature
R5	New feature added to allow users to keep multiple copies of photos	Inclusion of optional feature
R6	New feature added to send photo to other users by SMS	Inclusion of optional feature
R7	New feature added to store, play, and organize music. The management of photo (e.g. create, delete and label) was turned into an alternative feature. All extended functionalities (e.g. sorting, favorites and SMS transfer) were also provided	Changing of one mandatory feature into two alternatives
R8	New feature added to manage videos	Inclusion of alternative feature

Fig. 11. Summary of change scenarios in the MobileMedia SPL (based on [13, Tab.1])

2. As an SPL, it contains a large, rich set of varying (mandatory, optional and alternative) features which provides a wide set of representative aspects for study.
3. The original Java and AspectJ releases are available and were not written by us.

5.2 Change Propagation Analysis

A key benefit of a modular software design is in its ability to hide design decisions that are likely to change [27]. Thus, we consider the number of changed components as a result of a changed design decision to be an important comparator for a software design. To quantify this, similar to Figueiredo *et al.* [13], we measured the number of added, removed, and changed components in each system for each release.

Component Changes The changes to base components are shown in Figure 12. This table includes the pure Java releases (OO), pattern-based pointcut releases (PCD), annotation-based pointcut releases (@PCD), open modules releases (OM), and the quantified, typed event releases (EVT). This table considers Java classes and interfaces, aspects, and open modules.

Note that the declarations of annotations and event types are not included in the counts for this table, as they are measured separately and considered in the next section to give a direct comparison to pointcuts.

Components Added For all releases, new components added in the pattern-based pointcut (PCD) releases were also added to the annotation-based pointcut (@PCD), open

			R2	R3	R4	R5	R6	R7	R8	Total
Components	**Added**	OO	**9**	**1**	**0**	**5**	**7**	**10**	**6**	38
		OM	**17**	2	3	6	**11**	**17**	**22**	78
		PCD	13	2	3	6	8	14	16	62
		@PCD	13	2	3	6	8	14	16	62
		EVT	13	2	**2**	6	8	14	16	61
		Differences to PCD marked in BOLD blue								
	Removed	OO	**0**	0	0	0	0	1	**1**	2
		OM	1	0	0	0	0	1	0	2
		PCD	1	0	0	0	0	1	0	2
		@PCD	1	0	0	0	0	1	0	2
		EVT	1	0	0	0	0	1	0	2
	Changed	OO	5	**8**	**5**	**8**	**6**	**19**	**17**	68
		OM	5	**14**	**6**	**13**	**6**	**34**	**26**	104
		PCD	5	10	2	10	5	27	18	77
		@PCD	5	**8**	2	**11**	**7**	27	**20**	80
		EVT	5	**9**	**1**	**8**	5	**25**	**20**	73

Fig. 12. Base components change propagation in MobileMedia for each release

modules (OM), and quantified, typed event (EVT) releases. Note that the @PCD values are identical to the PCD values.

In R2, R6, R7, and R8, the number of added components differs for the open modules (OM) releases compared to PCD (marked in bold) due to the addition of modules in each of those releases. All aspects and base components in the OM releases are identical to the PCD releases.

In R4, the releases with pointcuts (PCD, @PCD, and OM) added an aspect that only handles precedence. This aspect was not added in the quantified, typed event release, as precedence in that release is controlled by the order of registering handler classes. This registration occurs inside the main class.

Components Removed In all 7 changed releases (R2–R8), the AO releases all have the same components removed. In R2, the PCD release removed a class `BaseThread` and in R8 the OO release removed the class `SplashScreen`. Since we did not implement either of the OO or PCD releases, we simply mimicked these changes in the @PCD, OM, and EVT releases.

Components Changed The difference between the components changed for the pattern-based pointcuts (PCD) and open modules (OM) releases is due entirely to changes in the modules, as once again all aspects and base components in the OM releases are identical to the PCD releases. Starting with R3, each release modified modules from the prior release due to changes in the aspects.

In R3, the PCD release changes two more components (`UtilAspectEH` and `ControllerAspectEH`) than the @PCD release due to the fragility of the pointcuts in those components. However, in R5, R6, and R8 the @PCD releases change more base components than the PCD releases, despite avoiding the fragile pointcut problem

with existing pointcuts. This is due to the need to annotate the base code with new annotations.

The difference in changes for R3 between @PCD and EVT was due to a changed event type requiring a change in the signature of the handler method.

For R4 however, the difference in values represents two important differences in the AO interfaces. First, the changed component in EVT was due to adding a precedence declaration to a handler (in @PCD this was a new aspect, not a changed aspect). Second, the two changed components in PCD and @PCD were from refactoring base code to expose join points. EVT did not need to perform such refactorings as it allows arbitrary statements as event announcements.

Of the remaining 7 changes that occurred in @PCD and not EVT, 3 were due to updating the precedence aspect, 1 was due to exposing join points and the remaining 3 were from changes in context (which for EVT shows up as changes in the event types).

Quantification Mechanism Changes The change propagation results are shown in Figure 13. The table lists the number of pointcuts added, changed, or removed for the open modules (OM), annotation-based pointcut (@PCD), and pattern-based pointcut (PCD) releases. The number of annotations added, changed or removed are shown for the @PCD releases and the number of event types added, changed, or removed are shown for the EVT releases.

Pointcuts		R2	R3	R4	R5	R6	R7	R8	Total
Add	OM	87	19	18	6	21	53	58	262
	PCD	64	12	13	4	15	39	43	190
	@PCD	64	12	13	4	15	39	43	190
Remove	*Differences to PCD marked in BOLD blue*								
	OM	0	0	0	0	2	12	11	25
	PCD	0	0	0	0	1	6	8	15
	@PCD	0	0	0	0	1	6	8	15
Change	OM	0	10	0	29	2	104	9	154
	PCD	0	9	0	18	2	74	4	107
	@PCD	0	4	0	13	2	65	4	88

Events/Anns		R2	R3	R4	R5	R6	R7	R8	Total
Add	@PCD	24	7	1	2	6	11	5	56
	EVT	16	4	0	2	6	5	3	36
Rem	*Differences to EVT marked in BOLD red*								
	@PCD	0	0	0	0	1	0	0	1
	EVT	0	0	0	0	0	0	0	0
Ch	@PCD	0	1	0	0	0	0	0	1
	EVT	0	2	0	1	0	12	1	16

Fig. 13. AO interfaces change propagation in MobileMedia for each release

Pointcuts The pointcuts added, removed, and changed were measured for all releases with pointcuts (PCD, @PCD, and OM) and there are two sets of comparisons to note. First, the OM releases have more pointcuts added and changed in almost every release (marked in bold) when compared to the PCD releases. This is due to the additional pointcuts contained in the module definitions.

The second comparison is between the PCD and @PCD releases. In three releases, the @PCD releases have fewer changed pointcuts. This occurred due to the gained stability from using the annotation-matching pointcut syntax. In total, the @PCD releases have almost 18% fewer changed pointcuts compared to the PCD releases.

Annotations and Events The annotations for the @PCD releases and the event types for the EVT releases are similar in that both mark join points in the base code for aspect code to advise. The change propagation of these two mechanisms is also similar. The differences between them (marked in bold) occur for several reasons.

The event types in EVT contain typed context declarations, while annotations do not contain any context. As such, when types in the base code (used as context) change, any event type referencing those types must also be updated. This is why the annotations have no changes in any @PCD release (the change in R3 was a renamed annotation) and the EVT releases have several changes.

In R2, the difference in the number of added event types and annotations is due to EVT's lack of quantification failure. For example, the @PCD release had to create an annotation to mark a join point for use in a `within` pointcut due to quantification failure. The EVT release was also able to re-use more event types than the @PCD release, saving the addition of 7 event types.

Pointcuts vs Annotations/Events In R7 a mandatory feature was turned into two alternative features, leading to changes in the base components which propagated to the event types and event handlers for the EVT release. 10 of the 12 resulting event type changes were due to the renaming of base components passed as context in those events types. Consider on the other hand the PCD release which required changing 38 of the 74 pointcuts due to the fragility of those pointcuts.

In R8, several new alternate features were added to the system. The EVT release was able to re-use several existing event types, leading to the addition of only 3 new event types. Similarly, the @PCD release only required the addition of 5 new annotations. The PCD release however required adding 43 new pointcuts to the system.

In general, note that the total number of added event types and annotations are 81% and 70% fewer, respectively, than the total number of added pointcuts for PCD releases. Also note that the total number of changed event types is 85% fewer than the total number of changed pointcuts in the PCD releases and 82% fewer than the total pointcuts changed in the @PCD releases.

Summary In summary, for some releases quantified, typed events showed an improved ability to withstand changes in components. In particular, for releases where significant refactoring in the base components took place, the EVT designs were able to reduce the impact of these changes in the base code from the handlers. Additionally,

- the total number of added event types and annotations are less than a third the number of pointcuts added in the PCD releases, showing that event types and annotations are re-used by multiple pointcuts,
- the total number of changed event types is 85% fewer than the total number of changed pointcuts in the PCD releases and 82% fewer than the total pointcuts changed in the @PCD releases,
- the @PCD releases have almost 18% fewer changed pointcuts compared to the PCD releases due to the lack of fragile pointcuts, and
- the EVT and @PCD releases were both able to efficiently re-use events/annotations leading to fewer additions in releases adding alternate features.

5.3 Software Engineering Metrics

As previously discussed, the main difference between most AO interfaces and quantified, typed events is that the dependency between components that announce events is explicitly stated using announce expressions that name event types. With most AO interfaces, this dependency is implicitly defined by the language semantics. Explicitly naming event types or annotations introduces coupling. The main goal of this section is to study the change in coupling between components. In order to perform this evaluation, we used a subset of the metrics suite proposed by Chidamber and Kemerer [5], Fenton and Pfleeger [11], and subsequently refined by Garcia *et al.* [13, 15].

		R2	R3	R4	R5	R6	R7	R8
CBC	**OO**	32	40	40	65	80	103	131
	OM	35	50	59	94	121	159	217
	PCD	35	50	59	94	121	159	217
	@PCD	82	106	122	161	200	255	332
	EVT	74	100	120	159	203	271	371
LCOO	**OO**	123	194	224	241	296	311	365
	OM	147	244	266	259	369	502	534
	PCD	147	244	266	259	369	502	534
	@PCD	147	244	266	259	369	502	534
	EVT	123	162	171	257	365	426	539

Fig. 14. Coupling and Cohesion for MobileMedia

Coupling   Coupling between components (CBC) [5] is a measurement of coupling. A component is coupled to another component if it accesses a field or calls a method on it. Figure 14 shows the results of our measurements.

The @PCD and EVT releases all have upwards of twice as much explicit coupling in the system compared to the PCD and OM releases. This is due to the explicit marking of join points (with annotations and event type announcements). However, realize that the added coupling is not coupling between aspects and base code but rather aspects to event types and base code to event types. Thus, this coupling only creates a maintenance

issue if an event type changes (such as in R7). This added coupling however is what allows the EVT releases to avoid other maintenance issues, such as pointcut fragility and limited access to context.

Cohesion Lack of cohesion in operations [5] (LCOO) is a measurement of cohesion of the classes in the system, based on how similar operations use attributes of the class. If methods of a class operate on the same attributes, the class is said to be cohesive and has a lower LCOO value. LCOO for all releases was measured and is shown in Figure 14. Note that the PCD, @PCD, and OM releases all have the same values due to having the same methods/fields in classes and ITDs/advice in aspects.

In general, quantified, typed events have more cohesion (indicated by lower LCOO) than the pointcut-based approaches. This is mostly due to the lack of needing to refactor the base code to expose join points to the aspect code. Such refactored code often only works on a small sub-set of the fields in the class, making the class less cohesive.

		R2	R3	R4	R5	R6	R7	R8
LOC	**OO**	1159	1314	1363	1555	2051	2523	3016
	OM	1337	1570	1700	1928	2474	3207	3999
	PCD	1276	1494	1613	1834	2364	3068	3806
	@PCD	1452	1723	1852	2094	2664	3461	4257
	EVT	1427	1669	1781	2050	2646	3398	4254
NOC	**OO**	24	25	25	30	37	46	51
	OM	31	33	36	42	53	69	91
	PCD	27	29	32	38	46	59	75
	@PCD	51	60	64	72	85	109	130
	EVT	47	53	56	64	78	96	115
NOA	**OO**	62	71	74	75	106	132	165
	OM	82	99	108	112	149	187	237
	PCD	62	72	76	77	110	139	177
	@PCD	62	72	76	77	110	139	177
	EVT	71	92	96	101	144	175	217
NOO	**OO**	124	140	143	160	200	239	271
	OM	143	169	179	197	247	308	369
	PCD	143	169	179	197	247	308	369
	@PCD	143	169	179	197	247	308	369
	EVT	142	167	177	196	245	302	378

Fig. 15. The measured size metrics for MobileMedia

Size Metrics Figure 15 shows the number of components (NOC) and total lines of code (LOC) for each release. The number of components includes classes and interfaces for all releases. For the PCD, @PCD, and OM releases it also includes aspects. For OM it includes modules, @PCD includes annotations and EVT includes event types.

Lines of code were measured using a tool[5] that ignores comment and whitespace lines. All other lines were included and every component from NOC was included.

---

[5] Retrieved from: `http://reasoning.com/downloads.html`

Number of operations (NOO) was measured as the total number of methods in classes, introduced methods in aspects, advice bodies in aspects and handler methods in event handlers. Number of attributes (NOA) was measured as the total number of fields in classes or aspects (including inter-type declared fields) and the number of context variables in quantified, typed events.

As one would expect from creating so many events and annotations, the lines of code and number of components is higher for both @PCD and EVT. The number of attributes is also higher for EVT due to counting event type context variables as attributes.

Summary   In summary, our results show the total explicit coupling is higher in the annotation-based pointcut and quantified, typed event releases due to the interface added between base components and aspects. The increased coupling is a trade-off for the stability gained by the interface between aspect and base code, as the previous section clearly demonstrates.

6   Case Study: Health Watcher

This section contains our second studied project, a web-based application called Health Watcher [16, 24, 34].

6.1   Health Watcher Overview

Health Watcher is an application for users to file health complaints. The system was initially developed in 2001 and has undergone 9 releases to add new features and fix previous bugs. The 10 releases and their descriptions are shown in Figure 16.

Release	Description
R1	Health Watcher base
R2	Applied Command pattern to remove dependency on Servlets
R3	Applied State pattern to keep from updating Complaints after they are closed
R4	Applied Observer pattern for calls to Update method
R5	Applied Adapter pattern to the distribution behavior
R6	Applied Abstract Factory pattern for repositories and data structures
R7	Applied Adapter pattern to remove dependency on Servlets
R8	Applied Abstract Factory pattern to generalize distribution types
R9	Added new functionality
R10	Improved exception handling

Fig. 16. Summary of change scenarios in Health Watcher

We chose to study Health Watcher for the following reasons.

 1. It was successfully used in several previous AO studies [16, 24, 34].

2. As a real-world system, there are 10 different releases available. This represents a real-world evolution of the application and provides a wide set of varying aspectual features.

3. The original Java and AspectJ releases are available and were not written by us.

6.2 Change Propagation Analysis

As stated in the previous case study, we consider the number of changed components as a result of a change in a design decision to be an important comparator for a software design. This section performs our analysis on Health Watcher.

Component Changes The changes to base components are shown in Figure 17. This table includes the Java releases (OO), pattern-based pointcut releases (PCD), annotation-based pointcut releases (@PCD), open modules releases (OM), and the quantified, typed event releases (EVT). This table considers Java classes/interfaces, aspects, and open modules.

			R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	Total
Components	Added	OO	88	4	12	2	3	4	4	4	12	5	138
		OM	106	12	16	4	0	4	4	2	12	6	166
		PCD	101	11	16	3	0	4	4	2	12	6	159
		@PCD	101	11	16	3	0	4	4	2	12	6	159
		EVT	100	10	16	3	0	4	4	2	12	6	157
		Differences to PCD marked in BOLD blue											
	Removed	OO	0	0	0	0	1	0	0	0	0	2	3
		OM	0	0	0	0	0	0	0	0	0	1	1
		PCD	0	0	0	0	0	0	0	0	0	1	1
		@PCD	0	0	0	0	0	0	0	0	0	1	1
		EVT	0	0	0	0	0	0	0	0	0	1	1
	Changed	OO	0	22	6	15	16	2	27	3	23	48	162
		OM	0	27	9	9	1	3	27	5	23	55	159
		PCD	0	25	8	7	1	2	27	3	22	52	147
		@PCD	0	26	8	29	1	2	27	3	23	55	174
		EVT	0	26	8	32	1	2	27	3	23	54	176

Fig. 17. Base components change propagation in Health Watcher for each release

Components Added For all releases, new components added in the pattern-based pointcut (PCD) releases were also added to the annotation-based pointcut (@PCD), open modules (OM), and quantified, typed event (EVT) releases. Note that the @PCD values are identical to the PCD values (as annotations are considered separately in Figure 18).

In R1, R2, and R4, the number of added components differs for the open modules (OM) releases compared to PCD (marked in bold) due to the addition of modules in

each of those releases. All aspects and base components in the OM releases are identical to the PCD releases.

In R1, an aspect that only contains a declare parents statement was not added in the EVT release. This statement failed to compile with the abc based intertype declarations implementation. Instead, we manually modified the base classes to add the Serializable interface to the 2 types. This particular aspect did not change in the PCD releases, thus our work-around did not cause problems in later EVT releases.

In R2, the releases with pointcuts (PCD, @PCD, and OM) added an aspect that only handles precedence. This aspect was not added in the quantified, typed event release, as precedence in that release is controlled by the order of registering handler classes. This registration occurs inside the main class or using annotations in the handler classes.

Components Removed   Similar to MobileMedia, in all 9 changed Health Watcher releases (R2–R10), the AO releases all have the same components removed.

Components Changed   Unlike MobileMedia where the difference between the components changed for the pattern-based pointcuts (PCD) and open modules (OM) releases was due entirely to changes in the modules, in Health Watcher some of the aspects also were modified in order to give anonymous pointcuts names (for the modules to reference).

Also unlike MobileMedia, the components changed for @PCD and EVT are more in Health Watcher for R4 than the PCD release due to needing to add annotations and event announcements in base code. This was because fewer base components changed in the PCD release but over 20 had to be modified to add annotations and event announcement.

Quantification Mechanism Changes   The change propagation results in terms of modularization techniques are shown in Figure 18. The table lists the number of pointcuts added, changed, or removed for the open modules (OM), annotation-based pointcut (@PCD), and pattern-based pointcut (PCD) releases. The number of annotations added, changed or removed are shown for the @PCD releases. It also lists the number of event types added, changed, or removed for EVT.

Pointcuts   Again, the pointcuts added, removed, and changed were measured for all releases with pointcuts (PCD, @PCD, and OM). Once again, the OM releases have more pointcuts added and changed compared to the PCD releases, as the module definitions also contain named pointcuts.

Unlike the MobileMedia case study, Health Watcher had relatively stable pointcuts. As such, the only benefit observed in the @PCD releases occurred in R2, where 3 fewer pointcuts were changed.

Annotations and Events   Unlike the MobileMedia case study, annotations and event types in Health Watcher perform roughly the same in all releases, with the exception of R3. In this release, there were several (4) events that had to be duplicated: once with a void return type and once with a non-void return type. This was due to the advice being applied to multiple methods (with differing return types). The PCD releases simply marked all methods with the same annotation and the aspect was able to advise them

Pointcuts		R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	Total
Add	OM	57	16	36	16	0	0	0	6	0	30	161
	PCD	28	11	12	10	0	0	0	6	0	20	87
	@PCD	28	11	12	10	0	0	0	6	0	20	87
Remove		Differences to PCD marked in BOLD blue										
	OM	0	0	0	0	0	0	0	6	6	0	12
	PCD	0	0	0	0	0	0	0	4	4	0	8
	@PCD	0	0	0	0	0	0	0	4	4	0	8
Change	OM	0	4	0	0	0	0	1	2	6	3	16
	PCD	0	4	0	0	0	0	1	0	5	3	13
	@PCD	0	1	0	0	0	0	1	0	5	3	10

Events/Anns		R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	Total
Add	@PCD	13	2	2	4	0	0	1	0	0	6	28
	EVT	14	4	9	5	0	0	0	0	0	10	42
Rem		Differences to EVT marked in BOLD red										
	@PCD	0	0	0	0	0	0	0	0	0	0	0
	EVT	0	0	0	0	0	0	0	0	0	0	0
Ch	@PCD	0	0	0	0	0	0	0	0	0	0	0
	EVT	0	1	0	0	0	0	0	0	0	0	1

Fig. 18. AO interfaces change propagation in Health Watcher for each release

all, without regard to the return type. This problem also accounts for the extra events in R1 and R2.

Pointcuts vs Annotations/Events In general, note that the total number of added event types and annotations are 52% and 68% fewer, respectively, than the total number of added pointcuts for PCD releases. This result is similar to the MobileMedia results.

Summary In summary, the Health Watcher case study showed similar results to the MobileMedia case study. The noticeable differences between the studies were due to the fact that the Health Watcher study tended to simply add new aspects and avoid changing existing aspects and base code as much as possible while the MobileMedia study made significant modifications (in order to change mandatory features into optional ones).

– the total number of changed event types and annotations is significantly lower than the total pointcuts changed in the PCD releases,
– the total number of added event types and annotations are 52% and 68% fewer, respectively, than the total number of added pointcuts for PCD releases,
– the aspects were fairly stable and there were few pointcuts changed in the PCD releases, unlike in MobileMedia, leading to very little benefit seen for the @PCD releases.

6.3 Software Engineering Metrics

Similar to the MobileMedia study, in this section we examine the coupling between components in the system. Again, we use a subset of the metrics suite proposed by Chidamber and Kemerer [5] and Fenton and Pfleeger [11] and subsequently refined by Garcia *et al.* [15].

		R1	R2	R3	R4	R5	R6	R7	R8	R9	R10
CBC	**OO**	281	318	352	365	369	375	417	424	571	586
	OM	262	300	326	327	327	330	373	374	486	502
	PCD	262	300	326	327	327	330	373	374	486	502
	@PCD	279	321	352	382	382	385	430	432	561	593
	EVT	333	379	437	476	476	479	521	525	649	688
LCOO	**OO**	791	802	519	568	624	604	604	604	779	827
	OM	764	779	588	599	599	599	599	597	730	809
	PCD	764	792	601	612	612	612	612	610	730	809
	@PCD	764	792	601	612	612	612	612	610	730	809
	EVT	767	803	612	633	633	633	633	630	745	804

Fig. 19. Coupling and Cohesion for Health Watcher

Coupling   The @PCD and EVT releases have upwards of 18–33% as much explicit coupling in the system compared to the PCD and OM releases. Again, this is due to the explicit marking of join points (with annotations and event type announcements).

Cohesion   Lack of cohesion in operations [5] (LCOO) for all releases was measured and is shown in Figure 19. Note that the PCD, @PCD and OM releases all have similar values due to having the same methods/fields in classes and ITDs/advice in aspects.

In this particular study, for the PCD releases base code was not refactored to expose join points as all pointcuts target existing methods or classes. Thus, the EVT releases do not have more cohesion than the pointcut-based approaches. This differs from the results shown for MobileMedia.

Size Metrics   Figure 20 shows the number of components (NOC) and total lines of code (LOC) for each release. As before, the number of components includes classes and interfaces for all releases. For PCD, @PCD and OM releases it also includes aspects. For OM it includes modules, @PCD includes annotations, and EVT includes event types.

Similar to the MobileMedia study, the lines of code and number of components is higher for both @PCD and EVT due to adding annotations and event types. The number of attributes is also higher for EVT due to counting event type context variables as attributes.

The results for the size metrics follow the same general trends as the MobileMedia study. Note however that in this study, due to the lower number of lines of code for the

		R1	R2	R3	R4	R5	R6	R7	R8	R9	R10
LOC	**OO**	5990	6371	6896	7048	7233	7296	7318	7357	8800	8697
	OM	5777	6202	6840	6944	6944	6994	7019	7027	7980	7937
	PCD	5712	6116	6718	6807	6807	6856	6881	6890	7853	7791
	@PCD	5851	6262	6891	7077	7077	7126	7161	7170	8167	8159
	EVT	5948	6369	7128	7371	7371	7422	7450	7464	8485	8551
NOC	**OO**	88	92	104	106	108	112	116	120	132	135
	OM	108	120	136	142	142	146	150	152	164	169
	PCD	103	114	130	135	135	139	143	145	157	162
	@PCD	116	129	147	155	155	159	164	166	178	190
	EVT	118	133	158	168	168	172	176	178	190	205
NOA	**OO**	187	215	218	221	223	225	227	228	248	256
	OM	205	236	252	256	256	258	260	260	278	296
	PCD	187	216	220	221	221	223	225	225	243	252
	@PCD	187	216	220	221	221	223	225	225	243	252
	EVT	199	233	305	311	311	313	315	315	333	347
NOO	**OO**	527	557	701	715	750	766	782	787	881	894
	OM	556	574	724	733	733	745	761	761	821	847
	PCD	556	574	724	733	733	745	761	761	821	847
	@PCD	556	574	724	733	733	745	761	761	821	847
	EVT	556	574	724	734	734	746	762	762	822	845

Fig. 20. The measured size metrics for Health Watcher

PCD releases compared to the OO releases, most aspect releases have fewer lines of code than the same OO release.

Summary   Once again, the annotation-based pointcut and quantified, typed event releases showed higher coupling than the PCD and OM releases. These results are similar to the MobileMedia study.

7   Key Observations

We observed several key benefits to the studied designs. These benefits are outlined in Figure 21 and described in detail in this section.

	Inter-type Declarations	Type-hierarchy Modification	Exception Softening	Quantification Support	Non-Fragile Pointcuts	Context Information
OM	Yes	Yes	Yes	*Limited*	Yes	*Limited*
PCD	Yes	Yes	Yes	*Limited*	*No*	*Limited*
@PCD	Yes	Yes	Yes	*Limited*	Yes	*Limited*
EVT	Yes	*No*	*No*	**Full**	Yes	**Full**

Fig. 21. Observations across all studied AO interfaces

7.1 Inter-Type Declarations

A static feature of AspectJ that allows adding fields/methods to other classes is inter-type declarations (ITDs) [20]. This feature was recently added to the Ptolemy language (with the same syntax as AspectJ) and thus available for all studied AO interfaces.

In MobileMedia, ITDs are used mostly for two purposes: to add additional data (fields) to existing types (and manipulate that new data) and to provide alternate implementations of features. ITDs first show up in R3 and are heavily used in later releases which contain alternate features. For example in R8, ITDs are defined in 12 out of the 22 aspects (54%).

In Health Watcher, ITDs are used to add methods for timestamping complaints, starting the remote server for RMI and to implement a singleton pattern. In the system, a total of 3 out of 26 aspects (11.5%) contain ITDs.

7.2 Declare Parents

Similar to ITDs, type hierarchies in the base components can be extended in a modular manner using AspectJ's *declare parents*. This feature seems well suited to help handle alternate features in a system, but was not heavily used by the current design of the MobileMedia product-line.

In MobileMedia, only R8 contains declare parents statements to extend two type hierarchies by adding a new super-class to the base components. For the PCD, @PCD, and OM releases these effects were modular. For the EVT releases, the base components had to be modified (due to a compiler bug) and these changes were non-modular, but not invasive.

In Health Watcher, declare parents statements appear in 6 out of 26 (23%) aspects. The statements are used in several places to place marker interfaces onto a set of types, which are then advised by the pointcut patterns. This was a useful pattern in this system and while Ptolemy supports declare parents statements, the lack of a pattern form of quantification meant that these marker interfaces were not useful in those releases.

7.3 Softened Exceptions

AspectJ also has the ability to soften exceptions thrown in the base components [20] using *declare soft* statements. This was used in MobileMedia to help modularize the exception handling feature in R2 and in Health Watcher for the exception handling feature in R1 and R10, allowing the base components to no longer declare they throw checked exceptions handled by the aspects. The releases for @PCD and OM also contained such modularizations.

Currently Ptolemy does not have any similar constructs and thus the base components must still declare that these checked exceptions are thrown. There are both pros and cons of these declarations. The con is that a programmer must write these additional annotations. On the positive side, having these annotations makes the features in EVT releases completely (un)pluggable. In the PCD-based releases, if a feature that provides exception softening is unplugged, the compilation of the base components fails.

7.4 Quantification Support

Quantified, typed events give the programmer the ability to add event announcement for any arbitrary statement in the base components. The pattern and annotation-based pointcut approaches can only advise join points available in the provided pointcut language, such as method executions or calls. This often results in what Sullivan *et al.* called quantification failure [39] and is caused by incompleteness in the language's event model. Quantification failure occurs when the event model does not implicitly announce some kinds of events and hence does not provide pointcut definitions that select such events [39].

In MobileMedia, we observed several instances of quantification failure. For example, in R2 the aspects needed to advise a *while* loop and similarly in R3 the aspects needed to advise a *for* loop. To accommodate this, all pointcut-based releases (PCD, @PCD, OM) refactor the base components, for example moving these loops into newly added methods. By R8, a total of 5 refactorings were made to expose join points. This accounts for approximately 5% of the advised join points. The EVT releases did not suffer from this problem and thus these refactorings were not necessary.

In Health Watcher, we observed a different form of quantification failure. However, this time the failure was in the EVT releases and related to the handling of design rules that encapsulate entire types. As previously mentioned, the PCD releases used declare parents statements to add marker interfaces to several types. The aspects then used pattern pointcuts to target all method executions in sub-types of that marker interface. This was used for things such as making all methods in a class synchronized. Figure 22 shows the implementation for this design rule in PCD, which uses the marker interface `SynchronizedClasses` on two types and around advice to wrap the execution of all methods in those types in a **synchronized** statement.

```
1 private interface SynchronizedClasses {}

3 declare parents: EmployeeRepositoryArray || ComplaintRepositoryArray
4                  implements SynchronizedClasses;

6 Object around(Object o): this(o) && execution(* SynchronizedClasses+.*(..)) {
7   synchronized(o) { return proceed(o); }
8 }
```

Fig. 22. Pattern-based pointcut version of a design rule to encapsulate 2 types and make all their methods synchronized

For the EVT releases, we had to manually track this design rule across the releases. This meant that if the types involved added new methods we would need to remember the design rule and ensure those new methods also announced the proper event. For the Health Watcher example, this maintenance scenario did not occur (as the types involved did not evolve across releases) but it is important to note that we still had to be aware of the design rules and check them in each release - something the PCD releases did not require.

7.5 Fragile Pointcut Problem

As mentioned by Figueiredo *et al.* [13], the pattern-based pointcut releases of Mobile-Media suffer from a fragile pointcut problem [13,30,37]. This could be observed in R7, where a mandatory feature PHOTO is generalized into two alternative features PHOTO or MUSIC. This required modifying many pointcuts previously relying on an implicit matching of signatures in the base components.

The renaming of the base components itself is not a problem in the EVT releases and in fact requires no modification of events or handlers; the handlers will match on the event type which remains unchanged. If the event type is renamed (for example, to remain consistently named to the base components) then all handlers and events for that event type must be updated accordingly. The key difference in these two scenarios is that in the PCD case, the developer must be aware of which pointcuts matched the given join point (which can be aided with tools such as AJDT) while in the EVT case, the compiler will specify type errors for every publisher and subscriber for that event type, eliminating fragile pointcuts.

Since @PCD releases are structurally similar to the EVT releases, they also benefited from a lack of fragile pointcuts. Similarly, the OM releases also benefited from a lack of fragile pointcuts.

Fragile pointcuts were observed in releases 3, 5, and 7. In total, 19 out of the 107 pointcuts changed (18%) across all releases were due to fragile pointcuts. This problem has already been demonstrated in small examples, however, its appearance in PCD releases of MobileMedia presents real evidence that it could affect maintenance of PCD systems. The ability of EVT, @PCD, and OM to mitigate these risks shows that such problems, when they occur in practice, can be solved using these different AO interfaces.

7.6 Access to Context Information

AspectJ provides means to access context information from advised join points [20]. The type of information available to advice however is limited by the language, such as the receiver object, method arguments, etc. In MobileMedia and Health Watcher, there were several instances where this lack of flexible availability to context added complexity to the system. For example, the exception handling aspects needed access to a field in the controller class being advised. Thus, the field needed marked public, a getter method added, or the aspect marked as privileged. Either way, the aspect becomes coupled to the interface of the advised class. This was a problem for all pointcut-based releases (PCD, @PCD, OM).

This was also a key difference between the @PCD and EVT releases. While the @PCD releases provided similar benefits in terms of preventing fragile pointcuts, in terms of context exposure the annotations were not a sufficiently expressive quantification mechanism when compared to EVT.

8 Comparing the Studies

The two studies shown in this paper represent over 400k lines of code in total. Mobile-Media is a software product line with 8 different releases. Health Watcher is a system

with 10 different versions, representing a single system evolving over time. Both systems make heavy use of aspects for modularization. In this section, we examine what some of the similarities and differences are between these two case studies.

Several similarities arose from the data collected on these two case studies.

– Both systems show an increase in explicit coupling for the @PCD and EVT releases. This is due to making previously implicit coupling (from pointcut patterns) into explicit coupling. This new explicit coupling however provides maintenance benefits such as non-fragile pointcuts and improved access to context information.
– Both systems showed problems accessing context information in the pointcut-based releases (PCD, @PCD, and OM). The EVT releases did not suffer from this problem.
– As one might expect, both systems showed additional overhead in terms of creating annotations and event types and explicitly marking the base code with those annotations or announcing those event types.
– Both systems showed a large number of pointcuts for the PCD releases were fragile, while the OM, @PCD, and EVT releases avoided this issue.

There were some differences between the two case studies as well.

– Health Watcher gave a compelling example of where an implicit style pattern actually avoids pointcut fragility and provides for easier maintenance (e.g., see Figure 22). Such examples did not show in the MobileMedia study but help provide evidence of potential maintenance issues that can arise when using an explicit quantification mechanism.
– MobileMedia had 5% of its pointcuts targeting points not available in the pointcut language (called quantification failure). This problem did not appear at all in Health Watcher.
– The aspects in Health Watcher were relatively stable compared to those in MobileMedia. This was due mostly to the types of changes occuring in the system. MobileMedia had a lot of renamed components as it evolved whereas Health Watcher was generally adding new features that did not impact existing code.

9 Discussion

It is important to note that our measurements of the open module releases are all based on the AspectJ-based implementation of open modules [26] (which to our knowledge is the only open modules implementation available). Thus, some of the measured differences we see are due not necessarily to open modules as an AO interface but instead due to this specific implementation.

For example, any aspect containing an inter-type declaration was required by the compiler to have a `friend` declaration in the module for the class(es) being extended. This declaration has the effect of exposing every join point in the class to the aspect and significantly affected our design of the modules.

Also, the use of named pointcuts in this implementation required copying the full pointcut signature to the module definition. This has the effect of requiring updating two

locations (the original pointcut definition and the module) if that signature changes. The pointcut signatures change any time the exposed context types changed or the pointcut is renamed. Any difference in the number of changed pointcuts between the PCD and OM releases of Figure 13 and Figure 18 are a result of this problem.

Specifically, for MobileMedia this was a large problem in two releases (5 and 7) as a number of base components were renamed. This renaming caused multiple pointcuts to change and that effect was duplicated in the module definitions. This problem does not necessarily manifest itself in Open Modules, as originally defined by Aldrich [1], but is an artifact of this specific implementation.

Earlier in the paper we claimed that the OM releases did not suffer from fragile pointcuts. While this was the case for our particular study, this specific implementation of open modules does allow writing code which would suffer from this problem. To avoid this issue however, two things must be done. First, the aspects and modules must use the named pointcuts defined by the base code. Second, the person maintaining the base code (or the module containing it) must ensure that any changes to that code are reflected in the pointcuts contained in that module.

While open modules does allow assignment of blame (to the module maintainer) when the pointcuts no longer match the intended set of points in the program, the @PCD and EVT approaches studied go a step further and provide compiler detection and reporting of this problem. Thus the @PCD and EVT approaches provide an automated solution (for all systems) to the problem while the general usage of this implementation of open modules does not.

10   Conclusion and Future Work

Finding a good separation of concerns is an important problem. It is vital for improving the reliability and evolution of software systems. New modularization techniques enable improved separation of concerns. Their invention and refinement is thus equally important for maintaining intellectual control on the growing complexity of software systems. Pattern-based [20] and annotation-based [23] pointcuts, open modules [1, 26], and quantified, typed events [30] are examples of such modularization mechanisms.

In this paper, we presented a rigorous evaluation of these AO interfaces on two already well-substantiated case studies [13, 34]. The results of our change propagation and analysis using standard design metrics [5, 11, 15] show that annotation-based pointcuts and quantified, typed events help limit the impact of change, at the cost of increased explicit coupling. This coupling however is generally not a problem as it is to interface-like entities (annotations and event types), not between base components and/or aspects.

- The annotation-based releases have 18% fewer changed pointcuts than the PCD releases, due to a lack of fragile pointcuts.
- The total number of changed event types in MobileMedia is 74% fewer than the total number of changed pointcuts in the pattern-based releases and 66% fewer than the total pointcuts changed in the annotation-based releases.

Despite the similarites, quantified, typed events have several benefits over annotation-based pointcuts.

- Event types are flexible and do not suffer from quantification failure.
- The uniform access to context information avoided the need to break encapsulation by exposing fields to make them available to the aspects.

The pattern- and annotation-based releases also showed benefit over quantified, typed events for certain design rules.

- For the quantified, typed event releases, we had to be aware of and manually maintain design rules related to encapsulating entire types (e.g. to make an entire class synchronized).
- The pattern- and annotation-based pointcut releases and open modules releases used pointcuts to automatically maintain such design rules.
- Such design rules show cases where patterns do not exhibit fragile pointcut behavior, as the pointcuts are expected to capture all methods in the advised types.

The results of our study point out that no language feature was immune to all problems. There appears to be a need for a future language design that incorporates several of the language features studied here. We hope that this study serves as a guide when designing these future languages.

In the future we plan to perform a net options value analysis [3, 38] to investigate the trade-off between the higher coupling observed in the annotation-based pointcut and quantified, typed event releases and the stability gained by providing interfaces between aspects and base code.

As different join point models provide different types of expressiveness, additional interesting future work might explore how these different join point models impact change. Similarly, all features in this study used static deployment. A lot of work has been done on dynamically deploying aspects [4, 7, 8, 17, 28, 29, 31, 32], which may provide additional interesting results in future studies.

Acknowledgments

References

1. Aldrich, J.: Open Modules: Modular reasoning about advice. In: ECOOP '05. pp. 144–168 (2005)
2. Avgustinov, P., Christensen, A.S., Hendren, L., Kuzins, S., Lhotak, J., Lhotak, O., de Moor, O., Sereni, D., Sittampalam, G., Tibble, J.: abc: an extensible AspectJ compiler. In: AOSD. pp. 87–98 (2005)
3. Baldwin, C.Y., Clark, K.B.: Design Rules, Vol. 1: The Power of Modularity. MIT Press (2000)

4. Bockisch, C., Haupt, M., Mezini, M., Ostermann, K.: Virtual machine support for dynamic join points. In: AOSD '04. pp. 83–92 (2004)
5. Chidamber, S.R., Kemerer, C.F.: A metrics suite for object oriented design. IEEE TSE 20(6), 476–493 (1994)
6. Dyer, R., Bagherzadeh, M., Rajan, H., Cai, Y.: A preliminary study of quantified, typed events. In: ESCOT '10 (2010)
7. Dyer, R., Rajan, H.: Nu: a dynamic aspect-oriented intermediate language model and virtual machine for flexible runtime adaptation. In: AOSD '08. pp. 191–202 (2008)
8. Dyer, R., Rajan, H.: Supporting dynamic aspect-oriented features. ACM Transactions on Software Engineering and Methodology (TOSEM) 20(2) (2010)
9. Dyer, R., Rajan, H., Cai, Y.: An exploratory study of the design impact of language features for aspect-oriented interfaces. In: AOSD '12. pp. 143–154 (2012)
10. Ekman, T., Hedin, G.: The JastAdd system — modular extensible compiler construction. Sci. Comput. Program. 69(1-3), 14–26 (2007)
11. Fenton, N.E., Pfleeger, S.L.: Software Metrics: A Rigorous and Practical Approach. Course Technology (1998)
12. Ferrari, F., Burrows, R., Lemos, O., Garcia, A., Figueiredo, E., Cacho, N., Lopes, F., Temudo, N., Silva, L., Soares, S., Rashid, A., Masiero, P., Batista, T., Maldonado, J.: An exploratory study of fault-proneness in evolving aspect-oriented programs. In: ICSE'10. pp. 65–74 (2010)
13. Figueiredo, E., Cacho, N., Sant'Anna, C., Monteiro, M., Kulesza, U., Garcia, A., Soares, S., Ferrari, F., Khan, S., Castor Filho, F., Dantas, F.: Evolving software product lines with aspects: an empirical study on design stability. In: ICSE (2008)
14. Filho, F.C., Cacho, N., Figueiredo, E., Maranhão, R., Garcia, A., Rubira, C.M.F.: Exceptions and aspects: The devil is in the details. In: FSE (2006)
15. Garcia, A., Sant'Anna, C., Figueiredo, E., Kulesza, U., Lucena, C., von Staa, A.: Modularizing design patterns with aspects: a quantitative study. In: AOSD. pp. 3–14 (2005)
16. Greenwood, P., Bartolomei, T.T., Figueiredo, E., Dósea, M., Garcia, A.F., Cacho, N., Sant'Anna, C., Soares, S., Borba, P., Kulesza, U., Rashid, A.: On the impact of aspectual decompositions on design stability: An empirical study. In: ECOOP. pp. 176–200 (2007)
17. Hirschfeld, R.: Aspect-oriented programming with AspectS. In: Net.Object Days '02 (2002)
18. Hoffman, K.J., Eugster, P.: Towards reusable components with aspects: an empirical study on modularity and obliviousness. In: 30th International Conference on Software Engineering (ICSE). pp. 91–100 (2008)
19. Inostroza, M., Tanter, E., Bodden, E.: Join point interfaces for modular reasoning in aspect-oriented programs. In: ESEC/FSE. pp. 508–511 (2011)
20. Kiczales, G., Hilsdale, E., Hugunin, J., Kersten, M., Palm, J., Griswold, W.G.: An overview of AspectJ. In: ECOOP (2001)
21. Kiczales, G., Lamping, J., Mendhekar, A., Maeda, C., Lopes, C., Loingtier, J.M., Irwin, J.: Aspect-oriented programming. In: ECOOP (1997)
22. Kiczales, G., Mezini, M.: Aspect-oriented programming and modular reasoning. In: 27th international conference on Software engineering (ICSE). pp. 49–58 (2005)
23. Kiczales, G., Mezini, M.: Separation of concerns with procedures, annotations, advice and pointcuts. In: ECOOP '05. pp. 195–213 (2005)
24. Kulesza, U., Sant'Anna, C., Garcia, A., Coelho, R., von Staa, A., Lucena, C.: Quantifying the effects of aspect-oriented programming: A maintenance study. In: International Conference on Software Maintenance (ICSM). pp. 223–233 (2006)
25. Masuhara, H., Kiczales, G.: Modeling crosscutting in aspect-oriented mechanisms. In: ECOOP. pp. 2–28 (2003)
26. Ongkingco, N., Avgustinov, P., Tibble, J., Hendren, L., de Moor, O., Sittampalam, G.: Adding Open Modules to AspectJ. In: AOSD '06. pp. 39–50 (2006)

27. Parnas, D.L.: On the criteria to be used in decomposing systems into modules. Communications of the ACM 15(12), 1053–8 (December 1972)
28. Popovici, A., Alonso, G., Gross, T.: Just-in-time aspects: efficient dynamic weaving for Java. In: AOSD '03. pp. 100–109 (2003)
29. Popovici, A., Gross, T., Alonso, G.: Dynamic weaving for aspect-oriented programming. In: AOSD '02. pp. 141–147 (2002)
30. Rajan, H., Leavens, G.T.: Ptolemy: A language with quantified, typed events. In: ECOOP (2008)
31. Rajan, H., Sullivan, K.J.: Eos: instance-level aspects for integrated system design. In: ESEC/FSE, pp. 297–306 (2003)
32. Rajan, H., Sullivan, K.J.: Classpects: unifying aspect- and object-oriented language design. In: ICSE, pp. 59–68 (2005)
33. Rashid, A., Moreira, A.: Domain models are not aspect free. In: MODELS '06 (2006)
34. Soares, S., Laureano, E., Borba, P.: Implementing distribution and persistence aspects with AspectJ. In: 17th conference on Object-oriented programming, systems, languages, and applications (OOPSLA). pp. 174–190 (2002)
35. Steimann, F.: Domain models are aspect free. In: MODELS/UML '05 (2005)
36. Steimann, F., Pawlitzki, T., Apel, S., Kastner, C.: Types and modularity for implicit invocation with implicit announcement. TOSEM '10 20(1) (2007)
37. Störzer, M., Koppen, C.: PCDiff: Attacking the fragile pointcut problem. In: European Interactive Workshop on Aspects in Software (September 2004)
38. Sullivan, K.J., Griswold, W.G., Cai, Y., Hallen, B.: The structure and value of modularity in software design. In: ESEC/FSE (2001)
39. Sullivan, K.J., Griswold, W.G., Rajan, H., Song, Y., Cai, Y., Shonle, M., Tewari, N.: Modular aspect-oriented design with XPIs. ACM TOSEM 20(2) (2009)
40. Young, T.: Using AspectJ to Build a Software Product Line for Mobile Devices. Master's thesis, UBC (2005)