

2011

Translucid contracts: Expressive specification and modular verification of aspect oriented interfaces

Mehdi Bagherzadeh
Iowa State University

Follow this and additional works at: <https://lib.dr.iastate.edu/etd>

 Part of the [Computer Sciences Commons](#)

Recommended Citation

Bagherzadeh, Mehdi, "Translucid contracts: Expressive specification and modular verification of aspect oriented interfaces" (2011).
Graduate Theses and Dissertations. 10400.
<https://lib.dr.iastate.edu/etd/10400>

This Thesis is brought to you for free and open access by the Iowa State University Capstones, Theses and Dissertations at Iowa State University Digital Repository. It has been accepted for inclusion in Graduate Theses and Dissertations by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Translucid contracts:
Expressive specification and modular verification of aspect oriented interfaces

by

Mehdi Bagherzadeh

A thesis submitted to the graduate faculty
in partial fulfillment of the requirements for the degree of
MASTER OF SCIENCE

Major: Computer Science

Program of Study Committee:
Hridesh Rajan, Major Professor
Vasant Honavar
Robyn R. Lutz
Samik Basu

Iowa State University

Ames, Iowa

2011

Copyright © Mehdi Bagherzadeh, 2011. All rights reserved.

DEDICATION

To my parents, Javad and Akram.

TABLE OF CONTENTS

LIST OF FIGURES	v
ACKNOWLEDGEMENTS	vii
CHAPTER 1. Introduction	1
1.1 Density of Join Point Shadows	1
1.2 Reasoning about Control Effects	3
1.3 Contributions	6
CHAPTER 2. Translucid Contracts	7
2.1 Program Syntax	7
2.2 Declarations	7
2.3 Expressions	8
2.4 Specification Features	9
CHAPTER 3. Verification of Programs with Translucid Contracts	11
3.1 Overview of Key Ideas in Verification	12
3.2 Checking Handler Refinement	13
3.3 Example Handler Refinement	14
3.4 Verifying Ptolemy Programs	16
3.4.1 Verification of Regular Methods	16
3.4.2 Verification of Handler Methods	18
3.4.3 Translation Function	18
3.4.4 Illustration of the Verification Algorithms	19
3.4.5 Runtime Assertion Checking (RAC)	20

CHAPTER 4. Analysis of Expressiveness	22
4.1 Direct Interference: Augmentation	22
4.2 Direct Interference: Narrowing	23
4.3 Direct Interference: Replacement	25
4.4 Direct Interference: Combination	26
4.5 More Expressive Control Flow Properties	27
CHAPTER 5. Applicability to Other AO Interfaces	30
5.1 Translucid Contracts for XPIs and AAIIs	30
5.2 Translucid Contracts for Open Modules	33
CHAPTER 6. Related Ideas	35
6.1 Contracts for Aspects	35
6.2 Modular Reasoning	36
6.3 Grey Box Specification and Verification	37
CHAPTER 7. Soundness of Reasoning	38
7.1 Substitution Algorithm	40
7.2 Proof of Soundness	43
CHAPTER 8. Conclusion and Future Work	47
BIBLIOGRAPHY	48

LIST OF FIGURES

Figure 1.1	A behavioral contract for aspect interfaces using Ptolemy [19] as the implementation language. See Section 2.1 for syntax.	2
Figure 1.2	A translucent contract for event type <code>Changed</code>	5
Figure 2.1	Ptolemy’s syntax [19], with refining expressions and contracts	8
Figure 2.2	Syntax for writing translucent contracts	9
Figure 3.1	Rules for checking structural refinement	14
Figure 3.2	Structural refinement relation (\sqsubseteq)	15
Figure 3.3	Handler refinement	15
Figure 3.4	Translation algorithm. The algorithm for converting program expressions into specification expressions that simulate running of handlers.	17
Figure 3.5	Translation of method <code>setX</code>	20
Figure 3.6	Unrolling translation function	20
Figure 3.7	Runtime assertion checking (RAC). Gray lines show pseudo code corresponding to generated code by the compiler.	21
Figure 4.1	Specifying augmentation with a translucent contract	22
Figure 4.2	After-augmentation handler	23
Figure 4.3	Specifying narrowing with a translucent contract	24
Figure 4.4	Narrowing handler	24
Figure 4.5	Specifying replacement with a translucent contract	25
Figure 4.6	Replacement handler	25
Figure 4.7	Combination contract and handler	27

Figure 4.8	Expressive control flow properties beyond [24]	28
Figure 5.1	Applying translucent contract to an XPI	30
Figure 5.2	Narrowing advice for XPI	31
Figure 5.3	Applying translucent contract to an AAI	32
Figure 5.4	Applying translucent contract to Open Modules	33
Figure 5.5	Narrowing handler for Open Module	34
Figure 7.1	Alternative operational semantics of Ptolemy[20]	40
Figure 7.2	Classes to simulate list of active objects	42
Figure 7.3	Substitution algorithm	43
Figure 7.4	Auxiliary functions of substitution algorithm	44
Figure 7.5	Structural similarity of translation and substitution of announce and invoke expressions	46

ACKNOWLEDGEMENTS

I like to take this opportunity to thank those who helped me through research and writing phases of this thesis. I like to thank Dr. Hridesh Rajan, my advisor, for the extraordinary role he has played in shaping my research interests. I also like to thank, Dr. Gary T. Leavens for the friendly cooperation he has always offered me. My thanks also goes to Sean L. Mooney which developed Ptolemy compiler, used to showcase the feasibility of the proposed ideas in this thesis. I also like to thank my committee members Dr. Vasant Honavar, Dr. Robyn R. Lutz and Dr. Samik Basu for their productive suggestions to improve the quality of this work.

Discussions with Dr. Gary T. Leavens, Dr. Eli Tilevich, Dr. Jonathan Aldrich and Dr. Kathryn McKinley, committee members of doctoral symposium of OOPSLA/SPLASH '10 conference (Systems, Programming, Languages, and Applications: Software for Humanity 2010), were of great help. Additionally, I like to thank the anonymous reviewers of AOSD '11 conference (Aspect Oriented Software Development 2011) and FOAL '10 workshop (Foundations of Aspect-Oriented Languages) for their detailed, insightful feedbacks about the current work. My thanks also goes to members of Laboratory for Software Design. Finally research work presented in this thesis has been made possible by the generous financial support of NSF (National Science Foundation) for the Ptolemy project.

CHAPTER 1. Introduction

Reasoning about aspect-oriented (AO) programs that use pointcuts and dynamic advice, as found in AspectJ programs, often seems difficult, due to two fundamental problems:

1. Join point shadows, i.e., places in the code where advice may apply, occur very frequently¹ And at each join point shadow, reasoning must take into account the effects of all applicable advice.
2. The control effects of advice must be understood in order to reason about a program's control flow and how advice might interfere with the execution of other advice.

1.1 Density of Join Point Shadows

As an example of the first problem, consider the straight-line code in below. In this listing, assuming that `x` and `y` are fields, there are at least 8 join point shadows, including the 5 method calls, the writes of `x` and `y`, and the read of `x`.

```
1 x = o1.m1(a.e1(), b.e2());
2 y = o2.m2(c.e3(), x);
```

Knowing what advice applies where is amenable to tool support. An example is the Eclipse AspectJ Development Tools (AJDT). The idea of aspect-aware interfaces [13], is equivalent to such tool support. However, the number of reasoning tasks grows with the number of join points and the amount of applicable advice.

One way of avoiding this problem of frequent occurrence of join point shadows, is to limit where advice may apply, for example, by using some form of explicit base-advice interface (AO interface), e.g. crosscutting interfaces (XPIs), open modules, etc, [1, 6, 19, 27, 28]. This is the approach we adopt

¹For example, a join point shadow occurs at each method or constructor call, and each field read and write.

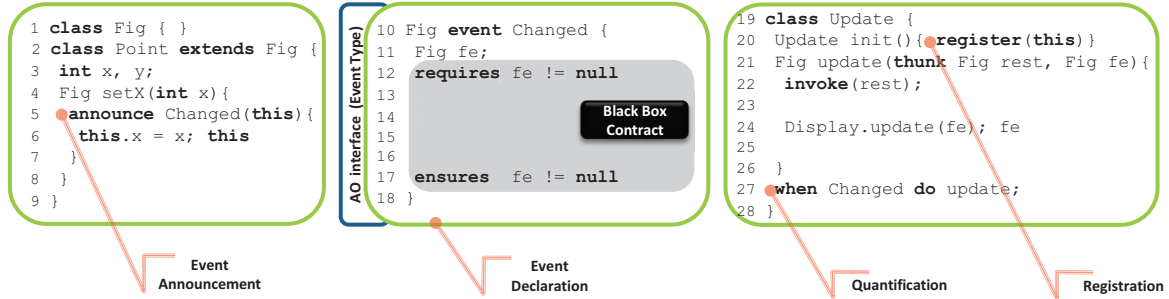


Figure 1.1 A behavioral contract for aspect interfaces using Ptolemy [19] as the implementation language. See Section 2.1 for syntax.

in this thesis by using the language Ptolemy [19]. Ptolemy introduces the notion of event types and limits the join points to explicit event announcements.

To illustrate, consider the Ptolemy code in Figure 1.1 from the canonical drawing editor example with functionalities to draw points, lines and update the display. In Ptolemy, events are explicitly announced, which mitigates the first problem, as reasoning about events only needs to happen at program points where events are explicitly announced (such as lines 5–7). Ptolemy programs declare event types, which are abstractions over concrete events in the program. Lines 10–18 declare an event type that is an abstraction over program events that cause change in a figure. An event type declaration may declare variables that make some context available. For example, on line 11, the changing figure, named fe , is made available. Concrete events of this type are *explicitly* and *declaratively* created using **announce** expressions as shown on lines 5–7. Like Eos [21, 22], Ptolemy doesn’t distinguish between aspects and classes. On lines 19–28 is the Ptolemy’s equivalent of an AspectJ-like advice, which advises calls to the method `setX`. The `Update` class has a binding declaration on line 27 that says to run the handler method `update` whenever events of type `Changed` are signaled. In Ptolemy’s terminology advice are called *handlers*. Ptolemy also provides dynamic registration using **register**, line 20, which activates the current instance of the `Update` class as an observer for the event `Changed`.

1.2 Reasoning about Control Effects

As an example of the second problem, understanding control effects of the advice, consider the `Logging` handler in the listing below that advises the same set of events advised by `Update` handler in Figure 1.1. To understand the control flow at these events matched by these handlers a developer must understand the control flow of both handlers. Furthermore, to understand the behavior at such events one must also understand the control flow of all other handlers that may advise the same events.

```

29 class Logging{
30 ...
31 Fig log(thunk Fig rest, Fig fe){
32   invoke(rest);
33   Log.logChanges(fe); fe
34 }
35 when Changed do log;
36 }

```

Design by contract (DBC) methodologies for aspect-oriented software development (AOSD) have been explored before [12, 28, 31], however, existing work relies on black box behavioral contracts. Such behavioral contracts specify, for each of the aspect’s advice methods, the relationships between its inputs and outputs, and treat the implementation of the aspect as a black box, hiding all the aspect’s internal states. As shown in Figure 1.1, event type `Changed` declares a black box contract on lines 12–17. Phrases “behavioral” and “black box” contract are used interchangeably throughout the thesis.

However, the black box contract on lines 12–17 does not specify the control effects of the handler.² For example, with just the black box contract of the event type `Changed` given, one cannot determine whether a call such as `p.setX(3)` will proceed to execute the body of `setX`, and thus whether such a call will always set the current `x` coordinate of `p` to its argument (3). If the expression `invoke` in the handler method `update` is forgotten inadvertently, the execution of the body of method `setX` will be skipped. This is equivalent to missing the call to `proceed` in an advice in AspectJ. Such assertions are important for reasoning, which depends on understanding the effect of composing the handler modules with the base code [24, 28]. That is, the contract does not specify if the handler must always proceed.

Ideas from Zhao and Rinard’s Pipa language [31], if applied to AO interfaces help to some extent. But, as discussed in Chapter 6, Pipa’s expressiveness beyond simple control flow properties is limited.

²This limitation of black box behavioral specifications was discussed in a preliminary version of this work [2].

Even if programmers don't use formal techniques to reason about their programs, contracts for AO interfaces can serve as the programming guidelines for imposing design rules [28]. But black box contracts for AO interfaces yield insufficiently specified design rules that leave too much room for interpretation, which may differ significantly from programmer to programmer. This may cause inadvertent inconsistencies in AO program designs and implementations, leading to hard to find errors.

Another problem with such black box contracts is that they do not help with effectively reasoning about the effects of aspects on each other. Consider another example concern, say `Logging`, which writes a log file at the events specified by `Changed`. For this concern different orders of composition with the `Update` concern in Figure 1.1 could lead to different results. (In AspectJ **declare precedence** can be used to enforce an ordering on aspects and the application of their advice.) Suppose line 22 of Figure 1.1 was omitted; that is, suppose that `Update` handler did not proceed. In that case, if `Update` were to run first, followed by `Logging`, then the evaluation of `Logging` would be skipped. Conversely, `Logging` would work (i.e., it would write the log file) if the handlers were composed in the opposite order. A handler developer cannot, by just looking at the black box contract of the event type, reason about the composition of such handlers. Rather a developer must be aware of the control effects of the code in all composed handlers. Furthermore, if any of these handlers changes (i.e., if their control effects change), one must reason about every other handler that applies at the same events.

The main contribution of this work is the notion of *translucid contracts* for AO interfaces, which is based on grey box specification [5]. A translucid contract for an AO interface can be thought of as an abstract algorithm describing the behavior of aspects that apply to that AO interface. The algorithm is abstract in the sense that it may suppress many actual implementation details, only specifying their effects using specification expressions. This allows the specifier to decide to hide some details, while revealing others. As in the refinement calculus, code satisfies an abstract algorithm specification if the code refines the specification [15], but we use a restricted form of refinement that requires structural similarity, to allow specification of control effects.

```

10 Fig event Changed {
11 ...
12 requires fe != null
13 assumes{
14   invoke(next);
15   establishes fe==old(fe)
16 }
17 ensures fe != null
18 }

```

Translucid Contract

```

19 class Update {
20 ...
21 Fig update(thunk Fig rest, Fig fe){
22   invoke(rest);
23   refining establishes fe==old(fe){
24     Display.update(fe); fe
25   }
26 }
27 ...
28 }

```

Figure 1.2 A translucid contract for event type Changed

We have added an example translucid contract to the AO interface, event type `Changed`, on lines 12–17 of Figure 1.2. Unlike a black box behavioral contract, internal states of the handler methods (which correspond to advice) that run when the event `Changed` is announced (this corresponds to a join point occurrence) are exposed in the translucid contract. In particular, any occurrence of the **invoke** expression (which is like AspectJ’s `proceed`) in the handler method *must* be made explicit in the translucid contract, line 14.³ This in turn allows the developer of the class `Point` that announces the event `Changed` to understand the control effects of the handler methods by just inspecting the specification of `Changed`. For example, from line 14 one may conclude that, irrespective of the concrete handler methods, the body for the method `setX` on line 6 of Figure 1.1 will always be run. Such conclusions allow a client of the `setX` to make more expressive assertions about its control flow without considering every handler method that may potentially run when the event `Changed` is announced. Expression **next** is a specification placeholder for the event closure passed to the handlers.

Requiring the **invoke** expression to be made explicit also benefits other handlers that may run when the event `Changed` is announced. For example, consider the logging concern discussed earlier. Since the contract of `Changed` describes the control flow effects of the handlers, reasoning about the composition of the handler method for logging and other handlers becomes possible without knowing

³**next** is a specification placeholder for the event closure passed to the handlers.

about all explicit handlers that may run when `Changed` is announced. In this thesis we explicitly focus on the use of translucent contracts for describing and reasoning about control flow effects.

To soundly reap these benefits, the translucent contract for the event type `Changed` must be refined by each conforming handler method [15]. We borrow the idea of structural refinement from JML’s model programs [25] and enhance it to support AO interfaces, which requires several adaptations that we discuss in Chapter 3. Briefly the handler method `update` on lines 22–25 in Figure 1.2 refines the contract on lines 12–17 because line 22 matches line 14 and lines 23–25 claim to refine the specification expression on line 15. The pre- and postconditions of `update` are considered the same as the pre- and postconditions of event type specification on lines 12 and 17, respectively.

1.3 Contributions

In summary, this work makes the following contributions:

- A specification and verification technique for writing contracts for AO interfaces and a proof of the soundness of the presented specification, verification and reasoning approach;
- An implementation of the proposed specification and verification technique in the Ptolemy’s compiler [18];
- An analysis of the effectiveness of our contracts using Rinard *et al.*’s work [24] on aspect classification which shows our technique works well for specifying all classes of aspects (as well as others that Rinard *et al.* do not classify);
- A demonstration that besides the AO interface proposal by the previous work of Rajan and Leavens [19], our technique works quite well for crosscutting interfaces [28] Aldrich’s open modules [1], and Kiczales and Mezini’s aspect-aware interface [13]. We also discuss the applicability of our technique to other languages that similarly solve the first reasoning problem by having explicit announcement, including Steimann *et al.*’s join point types [27], Hoffman and Eugster’s explicit join points [9]; and
- A comparison and contrast of our specification and verification approach with related ideas for AO contracts.

CHAPTER 2. Translucid Contracts

In this chapter, we describe our notion of translucid contracts and present a syntax to state these contracts. We use our previous work on the Ptolemy language [19] for this discussion.¹ However, as we show in Chapter 5 our basic ideas are applicable to other aspect-oriented programming models. We first present Ptolemy’s programming features and then describe its specification features.

2.1 Program Syntax

Ptolemy is an object-oriented (OO) language with support for declaring, announcing, and registering with events much like implicit-invocation (II) languages. The registration in Ptolemy is, however, much more powerful compared to II languages as it allows developers to quantify over all subjects that announce an event without actually naming them. This is similar to “quantification” in aspect-oriented languages such as AspectJ. The formally defined OO subset of Ptolemy has classes, objects, inheritance, and subtyping, but it does not have **super**, interfaces, exception handling, built-in value types, privacy modifiers, or abstract methods.

The syntax of Ptolemy executable programs is shown in Figure 2.1 and explained below. A Ptolemy program consists of zero or more declarations, and a “main” expression (see Figure 1.1 and Figure 1.2). Declarations are either class declarations or event type declarations.

2.2 Declarations

We do not allow nesting of *decls*. A class has a name (*c*) and names its superclass (*d*), and may declare fields (*field*) and methods (*meth*). Field declarations are written with a class name, giving

¹Descriptions of Ptolemy’s syntax and semantics are adapted from our previous work [19].

<pre> prog ::= $\overline{decl} e$ decl ::= class c extends d { $\overline{field} \overline{meth} \overline{binding}$ } t event p { $\overline{form} \overline{contract}$ } field ::= $t f$; meth ::= $t m$ (\overline{form}) { e } $t m$ (think $t var$, \overline{form}) { e } form ::= $t var$, where $var \neq \mathbf{this}$ and $var \neq \mathbf{next}$ binding ::= when p do m e ::= n var null new c () $e.m(\overline{e})$ $e.f$ $e.f = e$ $form = e$; e if (ep) { e } else { e } while (ep) { e } cast $c e$ e; e register (e) invoke (e) announce $p(\overline{e})$ { e } refining $spec$ { e } ep ::= n var $ep.f$ $ep \neq \mathbf{null}$ $ep == n$ $ep < n$! ep $ep \&\& ep$ </pre>	<p>where</p> <ul style="list-style-type: none"> $n \in \mathcal{N}$, the set of numeric, integer literals $c, d \in \mathcal{C}$, a set of class names $t \in \mathcal{C} \cup \{\mathbf{int}\}$, a set of types $p \in \mathcal{P}$, a set of event type names $f \in \mathcal{F}$, a set of field names $m \in \mathcal{M}$, a set of method names $var \in \{\mathbf{this}, \mathbf{next}\} \cup \mathcal{V}$, \mathcal{V} is a set of variable names
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 2.1 Ptolemy’s syntax [19], with **refining** expressions and contracts

the field’s type, followed by a field name. Method headers have a C++ or Java-like syntax, although their body is an expression. A binding declaration associates a set of events, described by an event type (p), to a method (m) [19]. An example is shown in Figure 1.2, which contains a binding on line 27. This binding declaration tells Ptolemy to run method `update` when events of type `Changed` are announced. II terminology calls such methods *handler methods*.

An event type (**event**) declaration has a return type (t), a name (p), zero or more context variable declarations (\overline{form}), and a translucent contract ($contract$). These context declarations specify the types and names of reflective information exposed by conforming events [19]. An example is given in Figure 1.2 on lines 10–18. In writing examples of event types, as in Figure 1.2, we show each formal parameter declaration ($form$) as terminated by a semicolon (;). In examples showing the declarations of methods and bindings, we use commas to separate each $form$.

2.3 Expressions

The formal definition of Ptolemy is given as an expression language [19]. It includes several standard object-oriented (OO) expressions and also some expressions that are specific to announcing events and registering handlers. The standard OO expressions include object construction (**new** c ()), variable dereference (var , including **this**), field dereference ($e.f$), **null**, cast (**cast** $t e$), assignment to a field ($e_1.f = e_2$), a definition block ($t var = e_1; e_2$), and sequencing ($e_1; e_2$). Their semantics and typing is fairly standard [6, 19] and we encourage the reader to consult [19].

There are also three expressions pertinent to events: **register**, **announce**, and **invoke**. The expression **register** (e) evaluates e to an object o , registers o by putting it into the list of active objects, and returns o . Only active objects in this list are capable of advising events. For example line 20 of Figure 1.2 is a method that, when called, will register the method’s receiver (**this**). The expression **announce** p (\bar{e}) $\{e\}$ declares the expression e as an event of type p and runs any handler methods of registered objects (i.e., those in the list of active objects) that are applicable to p [19]. The expression **invoke** (e) is similar to AspectJ’s **proceed**. It evaluates e , which must denote an event closure, and runs that event closure. This results in running the next handler method in the chain of applicable handlers in the event closure. If there are no remaining handler methods, it runs the original expression from the event. The type **think** t ensures that the value of the corresponding actual parameter is an event closure with return type t , and hence t is the type returned by **invoke**(e).

When called in an event, or by **invoke**, each handler method is called with a registered object as its receiver. The call passes an event closure as the first actual argument to the handler (`rest` in Figure 1.2 line 21). Event closures are never stored; they are only constructed by the semantics and passed to the handler methods.

There is one additional program expression: refining. A refining expression, of the form **refining** $spec$ $\{e\}$, is used to implement Ptolemy’s translucent contracts (see below). It executes the expression e , which is supposed to satisfy the contract $spec$.

2.4 Specification Features

The syntax for writing an event type’s contract in Ptolemy is shown in Figure 2.2. In this figure, all non-terminals that are used but not defined are the same as in Figure 2.1.

```

contract ::= requires  $sp$  assumes  $\{se\}$  ensures  $sp$ 
spec      ::= requires  $sp$  ensures  $sp$ 
 $sp$        ::=  $n$  |  $var$  |  $sp.f$  |  $sp \neq null$  |  $sp == n$  |  $sp < n$  |  $! sp$ 
            |  $sp == old(sp)$  |  $sp \&\& sp$ 
 $se$  ::=  $sp$  |  $spec$  | null | new  $c()$  |  $se.m(\bar{se})$  |  $se.f$  |  $se.f = se$  |  $form = se$ ;  $se$ 
        | if ( $sp$ )  $\{se\}$  else  $\{se\}$  | while ( $sp$ )  $\{se\}$  | cast  $c$   $se$  |  $se$ ;  $se$ 
        | register ( $se$ ) | invoke ( $se$ ) | announce  $p$  ( $\bar{se}$ )  $\{se\}$ 
        | refining  $spec$   $\{se\}$  | next | either  $\{se\}$  or  $\{se\}$ 

```

Figure 2.2 Syntax for writing translucent contracts

A *contract* is of the form **requires** sp_1 **assumes** { se } **ensures** sp_2 . Here, sp_1 and sp_2 are specification predicates as defined in Figure 2.2 and the body of the contract se is an expression that allows some extra specification-only constructs (such as the choice construct **either** se_T **or** se_F). In an event specification, the predicate sp_1 is the precondition for event announcement, and sp_2 is the postcondition of the event announcement. The specification expression se is the abstract algorithm describing conforming handler methods. The **invoke** expressions must be revealed in se and the variables that could be named in se are only context variables. If a method runs when an event of type p is announced, then its implementation must refine the contract se of the event type p . For example, in Figure 1.2, method `update`, lines 21–26 must refine the contract of the event `Changed`, lines 12–17.

There are four new expression forms that only appear in contracts: specification expressions, **next** expressions, abstract invoke expressions, and choice expressions. A specification expression (*spec*) hides implementation details (i.e., algorithms) and thus abstracts from a piece of code in a conforming implementation [23, 25]. The most general form of specification expression is **requires** sp_1 **ensures** sp_2 , where sp_1 is a precondition expression and sp_2 is a postcondition. Such a specification expression hides program details by specifying that a correct implementation contains a **refining** expression whose body expression, when started in a state that satisfies sp_1 , will terminate in a state that satisfies sp_2 [23, 25]. In examples we use the following syntactic sugars: **preserves** sp for **requires** sp **ensures** sp , and **establishes** sp for **requires** 1 **ensures** sp [23]. Ptolemy uses 0 for “false” and non-zero numbers, such as 1, for “true” in conditionals.

The **next** expression, the **invoke** expression and the choice expression (**either** – **or**) are placeholders in the specification that express the event closure passed to a handler, the call of an event handler using **invoke**, and a conditional expression in a conforming handler method, respectively. The choice expression hides the implementation details and thus abstracts from the concrete condition check in the handler method. For a choice expression **either** { se_1 } **or** { se_2 } a conforming handler may contain an expression e_1 that refines se_1 , or an expression e_2 that refines se_2 , or an expression **if** (e_0) { e_1 } **else** { e_2 }, where e_0 is a side-effect free expression, e_1 refines se_1 , and e_2 refines se_2 . Choice expression allows variability in handlers’ behaviors and enables their abstraction.

CHAPTER 3. Verification of Programs with Translucid Contracts

Verifying Ptolemy programs is different from standard object-oriented (OO) programs in two ways. First, a method in the program under verification may **announce** events that can cause a set of handlers to run. (In AspectJ, this is equivalent to invoking a set of advice at a join point.) Second, if the method is a handler it may call **invoke** that can also cause a set of handlers to run. (In AspectJ, this is equivalent to an advice calling **proceed** that can cause other advice to run.)

Therefore, verifying a Ptolemy program with translucid contracts poses two novel technical problems, compared to verifying standard OO programs: (1) verifying that each handler method correctly refines the contract of each event type it handles, and (2) verifying code containing **announce** and **invoke** expressions.

A *handler method* is a method that is statically declared in a *binding* form in its class to handle events of a given event type. When a *binding* of the form **when** p **do** m appears in a class declaration, then m is a *handler method for event type* p ; an example handler method is `update` in Figure 1.2.

The main novelty of translucid contracts is that both of these verification steps can be carried out modularly. By “modularly” we mean that each task can be done using only the code in question, the specifications of static types mentioned in the code, and the specifications of the relevant event types. For a handler, the relevant event types are all the event types that the method is a handler for (as determined by the binding declarations in the class where the handler is declared). For an **announce** expression, the relevant event type is the one that is being announced. For an **invoke** expression, which must occur inside a handler method body, it is each event type that the method is a handler for.

3.1 Overview of Key Ideas in Verification

Informally, to verify that each handler method correctly refines the contract of each event type that it handles, we first statically check whether the structure of the handler method body matches the structure of the **assumes** block of the event type. Note that **invoke** expressions that can override the underlying event body's execution (join point in AO terms) can only appear inside the handler method. So this check ensures that the control effects of the handler method matches the control effects specified in the translucent contract. At the same time, in our current implementation, we insert runtime assertions that check that the pre- and postconditions required by each event type's contract are satisfied by the handler method. These two checks ensure that starting with a state that satisfies the event type's precondition, if a correct handler method is run, it can only terminate in a state that satisfies the event type's postcondition, while ensuring that it produces no more control effects than those mentioned in the event type's **assumes** block.

Recall that an **announce** expression may cause a statically unknown number of handler methods to run, potentially followed by the event body. (In AspectJ terms, this is equivalent to running unknown number of pieces of advice, potentially followed by the original join point code.) An **invoke** expression (**proceed**) works similarly. To verify the code containing an **announce** expression, we take advantage of the fact that each correct handler method refines the event type's contract. So the event type's contract can be taken as a sound specification of the behavior of each handler. What is interesting and novel about our proposal is that the **assumes** block for an event type's translucent contract gives a sound specification of the behavior of an arbitrary number of handlers for that event.

Ignoring concrete details, imagine we need a sound specification of the behavior of the two handlers `Update` and `Logging` for the event type `Changed` in Figure 1.2. This can be constructed by taking the **assumes** block of this event type's contract and replacing occurrences of all **invoke** expressions inside it by the same **assumes** block (we will discuss how to do this shortly). This essentially achieves the effect of inlining the **invoke** expression (and is similar to unrolling a loop or inlining a recursive call [7]). Notice that construction of this specification only requires access to the event type. Also note that the resulting specification may contain some **invoke** expressions (as a result of inlining the **assumes** block). Let us call the constructed specification \mathcal{S} .

Given the specification \mathcal{S} of the behavior of the two handlers, we can now (1) reason about the code containing an **announce** expression as well as (2) the code containing an **invoke** expression. Again, ignoring concrete details, in the code containing the **announce** expression we do have access to the event body. So we replace all **invoke** expressions in \mathcal{S} with this event body. As a result, we now have a pure OO specification expression that is a sound specification of this announcement of the event `Changed`, \mathcal{S}_{ann} . This specification expression can be used to reason about the code that contains **announce** expression. An important property of this step is that we only used the event type’s contract and the code that was announcing events.

To reason about code that contains **invoke** expression, once again we start with a specification constructed from event type’s contract, e.g., \mathcal{S} . Note that the event body must refine the event type pre- and postcondition (to avoid surprising handler methods). So we replace all **invoke** expression in \mathcal{S} with the pre- and postcondition of the event type’s contract. This gives us a pure and sound OO specification of running two handlers and a correct event body, \mathcal{S}_{inv} . Similarly, in this step as well, we only used the event type’s contract and the code that contains **invoke** expression.

In the rest of this chapter, we describe these verification steps starting with the handler refinement.

3.2 Checking Handler Refinement

For sound modular reasoning, all handlers must be correct. A correct handler method in Ptolemy must refine the translucent contract of each event type that the method handles. Checking refinement of such a method is done in a two-step process. First, we statically verify whether the handler method’s body, which is an expression (e) is a structural refinement of the translucent contract of the event type, which is a specification expression (se). This step is performed as part of type-checking phase in Ptolemy’s compiler. Second, we verify that handler method satisfies the pre- and postconditions of the event type specification. This is currently checked at runtime (Section 3.4.5), however, a static approach, such as extended static checking [7], could also be applied.

Figure 3.1 shows the structural refinement process where refinement is checked for each handler method binding. CT is a fixed list of program’s declarations. Rule (CLASS TABLE REF) in Figure 3.1 checks structural refinement for each handler binding in the program. Rule (CHECK BINDING REF)

creates the typing contexts (π, Π) for the specification expression that is the body of the translucent contract and the program expression that is the body of the handler method and uses refinement rules in Figure 3.2 to check their structural refinement. In structural refinement, specification expressions in the contract are refined by program expressions in an implementation; however, program expressions in the contract are refined by textually identical program expressions in the refining implementation.

$$\begin{array}{c}
 \text{(CLASS TABLE REF)} \\
 \frac{\forall c \in \text{dom}(CT), \forall \text{binding} \in CT(c) \quad CT \vdash \text{binding in } c}{\vdash CT} \\
 \\
 \text{(CHECK BINDING REF)} \\
 \frac{\begin{array}{l}
 \text{decl} = \mathbf{t\ event} \ p \ \{t_1 \ \text{var}_1 \ \dots \ t_n \ \text{var}_n \ \text{contract}\}, \\
 \text{decl} \in CT, \quad \text{contract} = \mathbf{requires} \ sp_0 \ \mathbf{assumes} \ \{se\} \ \mathbf{ensures} \ sp_1, \\
 (t \ m(\mathbf{thunk} \ t' \ \text{var}'_0, t'_1 \ \text{var}'_1 \ \dots \ t'_m \ \text{var}'_m) \ \{e\}) \in CT(c), \quad \pi = \{\mathbf{next} : \mathbf{thunk} \ t, \text{var}_1 : t_1, \dots, \text{var}_n : t_n\}, \\
 \Pi = \{\mathbf{this} : c, \text{var}'_0 : \mathbf{thunk} \ t', \text{var}'_1 : t'_1, \dots, \text{var}'_m : t'_m\}, \quad (\pi, \Pi) \vdash se \sqsubseteq e
 \end{array}}{CT \vdash (\mathbf{when} \ p \ \mathbf{do} \ m) \ \text{in } c}
 \end{array}$$

Figure 3.1 Rules for checking structural refinement

A specification expression is refined by a program expression if its subexpressions are refined by corresponding subexpressions of the concrete program expression. Figure 3.2 shows key rules for checking that. There is no rule for **register** as it is not allowed in an event type specification. Judgement $(\pi, \Pi) \vdash se \sqsubseteq e$ states that specification expression se is refined by program expression e in the specification typing environment π and program expression typing environment Π , which in turn are constructed in the (CHECK BINDING REF) rule.

3.3 Example Handler Refinement

To illustrate the refinement rules in Figure 3.2, consider checking whether the handler method `update` on lines 22–25 in Figure 1.2 refines the translucent contract’s body on lines 14–15.

As illustrated in Figure 3.3 and according to the rule for $se_1; se_2$ in Figure 3.2, this refinement holds if (a) **invoke** (**next**) is refined by **invoke** (`rest`) and (b) **establishes** `fe==old(fe)` is refined by **refining establishes** `fe==old(fe) {Display.update(fe); fe}`.

For specification expression se , program expression e , specification and program typing contexts π and Π , se is refined by e , $(\pi, \Pi) \vdash se \sqsubseteq e$, as follows:		
Cases of Spec. Exp. (se)	Refined By (e)	Side Conditions
n	n	
var	var'	if $\pi(var) == \Pi(var')$
$sp.f$	$sp'.f$	if $(\pi, \Pi) \vdash sp \sqsubseteq sp'$
$sp! = null$	$sp'! = null$	if $(\pi, \Pi) \vdash sp \sqsubseteq sp'$
$!sp$	$!sp'$	if $(\pi, \Pi) \vdash sp \sqsubseteq sp'$
$sp_1 \&\& sp_2$	$sp'_1 \&\& sp'_2$	if $(\pi, \Pi) \vdash sp_1 \sqsubseteq sp'_1$, $(\pi, \Pi) \vdash sp_2 \sqsubseteq sp'_2$
$sp == n$	$sp' == n$	if $(\pi, \Pi) \vdash sp \sqsubseteq sp'$
$sp < n$	$sp' < n$	if $(\pi, \Pi) \vdash sp \sqsubseteq sp'$
$se_1; se_2$	$e_1; e_2$	if $(\pi, \Pi) \vdash se_1 \sqsubseteq e_1$, $(\pi, \Pi) \vdash se_2 \sqsubseteq e_2$
if (sp){ se_T } else { se_F }	if (ep){ e_T } else { e_F }	if $(\pi, \Pi) \vdash sp \sqsubseteq ep$, $(\pi, \Pi) \vdash se_T \sqsubseteq e_T$, $(\pi, \Pi) \vdash se_F \sqsubseteq e_F$
while (sp){ se }	while (ep){ e }	if $(\pi, \Pi) \vdash sp \sqsubseteq ep$, $(\pi, \Pi) \vdash se \sqsubseteq e$
$t \text{ var} = se_1; se_2$	$t \text{ var} = e_1; e_2$	if $(\pi, \Pi) \vdash se_1 \sqsubseteq e_1$, $\pi' = \pi \uplus \{var : (t, l)\}$, $\Pi' = \Pi \uplus \{var' : (t, l)\}$, $(\pi', \Pi') \vdash se_2 \sqsubseteq e_2$
refining spec { se }	refining spec { e }	if $(\pi, \Pi) \vdash se \sqsubseteq e$
$spec$	refining spec { e }	
invoke (se)	invoke (e)	if $(\pi, \Pi) \vdash se \sqsubseteq e$
announce $p(\bar{se})$ { se }	announce $p(\bar{e})$ { e }	if $(\pi, \Pi) \vdash \bar{se} \sqsubseteq \bar{e}$, $(\pi, \Pi) \vdash se \sqsubseteq e$
either { se_T } or { se_F }	if (ep){ e_T } else { e_F }	if $(\pi, \Pi) \vdash se_T \sqsubseteq e_T$, $(\pi, \Pi) \vdash se_F \sqsubseteq e_F$
either { se_T } or { se_F }	e_T	if $(\pi, \Pi) \vdash se_T \sqsubseteq e_T$
either { se_T } or { se_F }	e_F	if $(\pi, \Pi) \vdash se_F \sqsubseteq e_F$


Figure 3.2 Structural refinement relation (\sqsubseteq)

Figure 3.3 Handler refinement

For proving condition (a), we must check whether the subexpression **next** is refined by the subexpression **rest**. This can be done by the rule for var , which states that both variables **next** and **rest** must be given the same type by their respective typing contexts (π and Π). The specification typing

context π in this case, gives type **thunk** Fig to **next**, which is the same as the type for `rest` given by the program typing context Π . By applying the rule for *spec* in Figure 3.2, we can prove (b) because specification predicates **refining establishes** $\text{fe} == \text{old}(\text{fe})$ are the same in both specification expression and the program expression. Thus, the handler method `update` correctly refines the translucent contract for the event type `Changed`.

The refinement rule for the case *spec* deserves further explanation. It states that a specification expression *spec* is refined by an expression **refining** *spec* $\{e\}$, which claims to refine the same specification *spec*. The claim that *e* satisfies *spec* is discharged using runtime assertion checking as discussed in Section 3.4.5. The rules in Ptolemy’s operational semantics which discharge this condition are shown in Figure 7.1, rule (REFINING).

3.4 Verifying Ptolemy Programs

The main difficulty in verifying Ptolemy programs is that **announce** and **invoke** expressions could cause a statically unknown set of handlers (advice) to run. This set is not known statically unless a whole program analysis is performed. Thus such knowledge is not part of modular verification. Despite this, translucent contracts make modular verification possible. The challenge is to verify the code containing **announce** and **invoke** expressions. The basic idea is to use the translucent contract of the event type in place of each handler as discussed in Section 3.1. There are two types of methods to verify, regular methods which might announce event and handler method which handle the events.

3.4.1 Verification of Regular Methods

To statically verify a non-handler method $t\ m\ (\bar{t}\ \overline{\text{var}})\{e\}$ we must replace any occurrence of **announce** expression in its body *e* with a simulating expression for verification. The translation function *Tr* given in Figure 3.4 shows how to do that. Basically, a translation function $Tr(se, b_e, p)$ inlines event type specification/event body in place of announce/invoke expressions in *se*, as informally discussed in Section 3.1, to compute a simulating specification expression, modeling event announcement. Event *p* is the announced event, if any, and b_e is the event body. Function *Tr* is discussed in greater detail in Section 3.4.3.

For specification expressions se , expressions b_e , event types p , where p has contract requires sp_p assumes $\{se_p\}$ ensures sp'_p and context variables $\overline{t\ var}$, $Tr(se, b_e, p) =$		
Cases of se	Result	Side Conditions
$n, \mathbf{var}, \mathbf{null}, \mathbf{new}\ c(), \mathbf{next}, \mathit{spec}$	se	
$\mathbf{old}\ (se_1)$	$\mathbf{old}\ (se_2)$	if $se_2 = Tr(se_1, b_e, p)$
$se_1.f$	$se_2.f$	if $se_2 = Tr(se_1, b_e, p)$
$\mathbf{either}\ \{se_0\}\ \mathbf{or}\ \{se_1\}$	$\mathbf{either}\ \{se'_0\}\ \mathbf{or}\ \{se'_1\}$	if $se'_0 = Tr(se_0, b_e, p),$ $se'_1 = Tr(se_1, b_e, p)$
$se.m(\overline{se})$	$se'.m(\overline{se}')$	if $se' = Tr(se, b_e, p),$ $\overline{se}' = Tr(\overline{se}, b_e, p)$
$se_0.f = se_1$	$se'_0.f = se'_1$	if $se'_0 = Tr(se_0, b_e, p),$ $se'_1 = Tr(se_1, b_e, p)$
$\mathbf{if}(ep)\{se_0\}\ \mathbf{else}\{se_1\}$	$\mathbf{if}(ep')\{se'_0\}\ \mathbf{else}\{se'_1\}$	if $ep' = Tr(ep, b_e, p),$ $se'_0 = Tr(se_0, b_e, p),$ $se'_1 = Tr(se_1, b_e, p)$
$\mathbf{while}(ep)\{se\}$	$\mathbf{while}(ep')\{se'_0\}$	if $ep' = Tr(ep, b_e, p),$ $se' = Tr(se, b_e, p)$
$\mathbf{cast}\ c\ se$	$\mathbf{cast}\ c\ se'$	if $se' = Tr(se, b_e, p)$
$se_0; se_1$	$se'_0; se'_1$	if $se'_0 = Tr(se_0, b_e, p),$ $se'_1 = Tr(se_1, b_e, p)$
$\overline{t\ var} = se_0; se_1$	$\overline{t\ var} = se'_0; se'_1$	if $se'_0 = Tr(se_0, b_e, p),$ $se'_1 = Tr(se_1, b_e, p)$
$\mathbf{refining}\ \mathit{spec}\ \{se_1\}$	spec	
$\mathbf{register}(se_1)$	se_2	if $se_2 = Tr(se_1, b_e, p)$
$\mathbf{invoke}(se_1)$	$\mathbf{refining}\ \mathit{spec}\{\mathbf{either}\ \{se_2; b_e\}\ \mathbf{or}\ \{se_2; se_3\}\}$	if $se_2 = Tr(se_1, b_e, p),$ $se_3 = Tr(se_p, b_e, p),$ $\mathit{spec} = \mathbf{requires}\ sp_p$ ensures sp'_p
$\mathbf{announce}\ p'(\overline{se})\ \{se_1\}$	$\mathbf{refining}\ \mathit{spec}\{\mathbf{either}\ \{se'; se'_1\}\ \mathbf{or}\ \{\overline{t\ var}' = se'; se'_2\}\}$	if p' has translucent contract requires $sp_{p'}$ assumes $\{se_{p'}\}$ ensures $sp'_{p'}$ and context variables $\overline{t' \ var}'$, $\overline{se}' = Tr(\overline{se}, b_e, p),$ $se'_1 = Tr(se_1, b_e, p'),$ $se'_2 = Tr(se_{p'}, se'_1, p'),$ $\mathit{spec} = \mathbf{requires}\ sp'_{p'}$ ensures $sp'_{p'}$

Figure 3.4 Translation algorithm. The algorithm for converting program expressions into specification expressions that simulate running of handlers.

For the method m above with the body of e , we compute $Tr(e, \mathbf{skip}, \perp)$. The arguments **skip** and \perp specify that this method does not handle any events (\perp) and thus there is no event body (**skip**) which basically means the method is a non-handler. These parameters are included in this case simply to facilitate uniform application of the Tr function for both regular (non-handler) and handler methods.

The result of $Tr(e, \mathbf{skip}, \perp)$ is a specification expression with no Ptolemy-specific features, but may have extra expressions which simulate event announcement and running of handlers. This expression can then be used to perform standard weakest precondition based verification for OO programs.

3.4.2 Verification of Handler Methods

To statically verify a handler method h of the form $t\ h\ (\mathbf{thunk}\ t_0\ var_0, \bar{t}\ \overline{var})\ \{e\}$, for each event type p with a binding $\mathbf{when}\ p\ \mathbf{do}\ h$, one does the following. Let the contract for p be $\mathbf{requires}\ sp_p\ \mathbf{assumes}\ \{se_p\}\ \mathbf{ensures}\ sp'_p$, then compute $Tr(e, \mathbf{requires}\ sp_p\ \mathbf{ensures}\ sp'_p, p)$ and use the result to verify the handler h . The second argument to Tr is a specification statement consisting of the event's pre- and postconditions; this is used in the place of the announced event's body, since the event body is not available during static verification of the handler, and since this specification statement must be refined by all event bodies. The result of $Tr(e, \mathbf{requires}\ sp_p\ \mathbf{ensures}\ sp'_p, p)$ is a pure OO specification expression.

3.4.3 Translation Function

As illustrated in Section 3.1, the translation function $Tr(se, b_e, p)$, with p as the announced event and b_e as the body of p , inlines event type specification or event body in the place of announce and invoke expressions in se , and computes a simulating specification of the event announcement. Announce and invoke expressions are replaced by the event type's contract if there are more applicable handlers and are replaced by the event body otherwise. As existence or non-existence of more applicable handlers is not decidable statically, the translation algorithm considers occurrence of both of these situations simultaneously using an **either** – **or** choice expression, as shown in Figure 3.4.

Most cases in the translation function Tr are straightforward as they just recursively apply Tr to their subexpressions and compose the results. Translations of refining, announce and invoke expressions are of more interest, though. Translation of **refining** $spec\ \{e\}$ is $spec$ as the runtime assertion checking ensures that e refines the $spec$. The cases for invoke and announce expressions are central as they model event announcement by simulating running of the handlers and the event body.

Translations of invoke and announce expressions, both produce an **either** – **or** choice expression guarded by a **refining** expression. The either-branch simulates the situation when there is no applicable handler whereas the or-branch handles the situation when there exist more handlers to run.

In the translation of **invoke** (se_1) , the either-branch contains a sequence of two expressions: translation of the argument se_1 and the event body b_e , which means no more handler to run. The or-

branch contains a sequence of two expressions too: translation of argument se_1 and translation of the translucent contract se_p . The guarding refining expression assures that specification $spec$ is satisfied by the choice expression inside. $spec$ contains pre- and postcondition of the contract se_p .

Translation of announce expression is similar to the invoke. In case of **announce** $p' (\overline{se}) \{se_1\}$, the either-branch contains a sequence of two expressions: translation of the argument \overline{se} and the translation of event body se_1 . In or-branch the first expression is the translation of the arguments and their assignment them to context variables $\overline{var'}$. The second expression is the translation of the translucent contract of event p' , i.e. $se_{p'}$, assuming that the event body is se'_1 , the translation of se_1 . The translation of $se_{p'}$ simulates running of handlers for event p' with a concrete event body and event type's translucent contract as an abstraction for handlers.

The translation function $Tr(e, b_e, p)$ treats e as a subset of $se \cup \{spec\}$. Since the syntactic set $se \cup \{spec\}$ is a strict superset of syntactic set e , for every expression e there is an equivalent expression in the set $se \cup \{spec\}$.

The translation function assumes an acyclic event announce/handle relation. Circular relations could simply be detected statically.

3.4.4 Illustration of the Verification Algorithms

To illustrate, consider verifying the method `setX` in Figure 1.1 with the translucent contract in Figure 1.2. The body of this method is the announce expression **announce** `Changed(this) { this.x = x; this }`. To verify this method, we first apply the translation function $Tr(se, \mathbf{skip}, \perp)$ with $se = \mathbf{announce} \text{ Changed}(\mathbf{this})\{\mathbf{this}.x = x; \mathbf{this}\}$ as this method is a non-handler regular method. The case for announce expression in Figure 3.4 is applicable, which results in the specification expression shown in Figure 3.5.

Notice the use of the translation function Tr on lines 4–5. To verify this expression both the either-branch and the or-branch must be verified. During the verification, upon reaching the translation function, it is unrolled one more time resulting in the specification expression shown in Figure 3.6.

During this application, the cases for sequence, $spec$ and invoke expressions are used, which again results in an embedded translation function Tr on lines 6–7. The astute readers may have observed

```

1 refining requires fe != null ensures fe!= null{
2   either { this ; this.x = x; this }
3   or { Fig fe = this ;
4       Tr(invoke(next); establishes fe==old(fe),
5         this.x = x; this, Changed) }
6 }

```

Translation function

Figure 3.5 Translation of method setX

```

1 refining requires fe!= null ensures fe!= null{
2   either { this ; this.x = x; this }
3   or { Fig fe = this ;
4       refining requires fe!= null ensures fe!= null{
5         either { next; this.x = x; this }
6         or { next; Tr(invoke(next); establishes fe==old(fe),
7           this.x = x; this, Changed) }
8       }
9       establishes fe == old(fe) }
10 }

```

Unrolling translation function

Figure 3.6 Unrolling translation function

that we have essentially reduced problem of verifying **announce** and **invoke** expressions to a problem similar to reasoning about loops. Thus, standard techniques for reasoning about loops, such as proof rules that rely on user-supplied invariants, could be applied here. Heuristics like the one used in ESC/Java [7] to unroll the loops are also applicable here. When the verifier decides to terminate recursive unrolling, based on any of the above-mentioned approaches, the translation function in the result expression is just ignored. Verification of the method `update` is similar.

3.4.5 Runtime Assertion Checking (RAC)

As previously mentioned, some of the verification obligations encountered during the verification are discharged by relying on runtime assertions. Runtime checking discharges the following obligations, verifying that: (1) each handler method satisfies the specification of the event types it handles (2) each event body satisfies the pre- and postconditions of its event type specification, (3) each **refining** expression body refines the specification it claims to refine, and (4) each event announce-

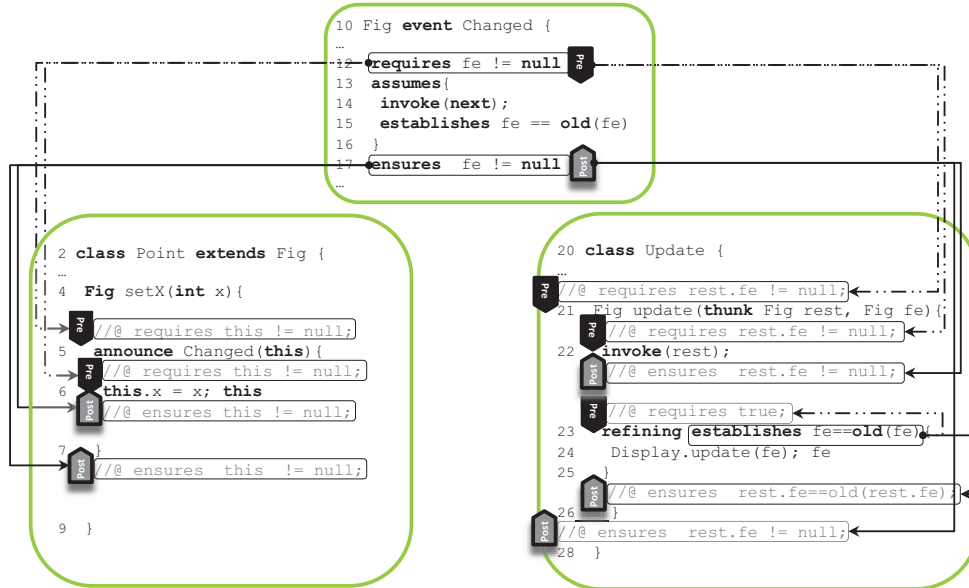


Figure 3.7 Runtime assertion checking (RAC). Gray lines show pseudo code corresponding to generated code by the compiler.

ment and consequent execution of all of its handler methods combined behavior, satisfies pre- and postconditions of the event type, regardless of the number of the handlers and their order of execution. Alternatively, a static checker like ESC/Java [7] could discharge these assumptions.

We have implemented runtime assertion checking in the Ptolemy compiler [18]. Figure 3.7 illustrates insertion of runtime probes by the Ptolemy compiler in the generated code. An abstraction function matches up context variable `fe` to its corresponding variables in the scopes of subject `Point` and handler `Update`.

To meet obligation (1) pre- and postcondition probes are inserted at the beginning and end of handler method body, before line 21 and after line 26. Runtime probes right before and after line 6 guarantee obligation (2). To verify that the refining expression on lines 23-25 refines the specification it claims to refine, obligation (3), runtime assertions are inserted before line 23 and after line 25. Finally to assure obligation (4) that event announcement and execution of handler methods does not violate the event type pre- and postconditions, runtime checks are enforced before and after **announce** and **invoke** expressions in the code. Runtime probes before line 5 and after line 7 guarantee the obligation for **announce** whereas probes right before and after line 22 meet the obligations for **invoke**.

CHAPTER 4. Analysis of Expressiveness

To analyze the expressiveness of translucent contracts, in this section we illustrate their application to specify base-aspect interaction patterns discussed by Rinard *et al.* [24]. Rinard *et al.* classify base-advice interaction patterns into: *direct* and *indirect interference*. Direct interference is concerned about control flow interactions whereas indirect interference refers to data flow interactions. Direct interference is concerned about calls to **invoke**, which is the Ptolemy's equivalent of AspectJ's **proceed**. Direct interference is further categorized into 4 classes of: augmentation, narrowing, replacement and combination advice which call **invoke** exactly once, at most once, zero and any number of times, respectively. An example, built upon the drawing editor example in Chapter 1, is shown for each category of the direct interference.

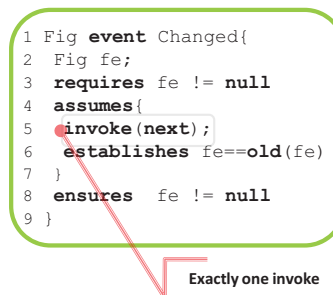
4.1 Direct Interference: Augmentation

Informally an augmentation handler evaluates **invoke** expression exactly once. An augmentation handler can be a before or after handler. After-augmentation handler is executed after the event body whereas in the before augmentation the order is opposite.

```

1 Fig event Changed{
2   Fig fe;
3   requires fe != null
4   assumes{
5     invoke(next);
6     establishes fe==old(fe)
7   }
8   ensures fe != null
9 }

```



Exactly one invoke

Figure 4.1 Specifying augmentation with a translucent contract

To illustrate consider the translucent contract in Figure 4.1 on lines 3–8. Translucent contracts are required to reveal all appearances of the `invoke` expression, thus it is assured that all refining handlers will evaluate `invoke` expression exactly once.

Furthermore, `invoke` is called at the beginning of the contract, requiring event handlers to run after the event body which means not only the refining handlers are augmentation handlers, but also that they run after the event body, after-augmentation handlers.

Method `log` in class `Logging` in Figure 4.2 is an example of a conforming after-augmentation handler. The requirement for this method is “to log the changes when figures are changed”. The handler `log` causes the event body to be run first by calling `invoke` on line 12 and then logs the changes in the figure on line 14. The classes `Point` and `Fig` are the same as in Figure 1.1.

```

10 class Logging{
11   Fig log(thunk Fig rest, Fig fe){
12     invoke(rest);
13     refining establishes fe==old(fe){
14       Log.logChanges(fe); fe
15     }
16   }
17   when Changed do log;
18 }

```

Figure 4.2 After-augmentation handler

Structural similarity requires the handler implementation to evaluate `invoke` exactly once and at its very beginning which in turn ensures that the handlers is an “after-augmentation” handler. The handler refines the contract because line 12 matches line 5 and the refining expression on lines 13–15 refines the same specification as on line 6.

4.2 Direct Interference: Narrowing

A narrowing handler evaluates `invoke` at most once, which implies existence of a conditional statement guarding `invoke`.

To illustrate consider the translucent contract in Figure 4.3 on lines 5–8 which specifies narrowing handlers. The contract reveals appearances of `invoke` expression and the `if` expression guarding that which in turn ensures that `invoke` expression is evaluated at most once. It does not, however, reveal

```

1 Fig event Changed{
2   Fig fe;
3   requires fe != null
4   assumes{
5     if(fe.fixed == 0)
6       invoke(next)
7   else
8     establishes fe==old(fe)
9   }
10  ensures fe != null
11 }

```

At most one invoke

Figure 4.3 Specifying narrowing with a translucent contract

the actual code that must refine the specification on line 8. All the refining handlers will have the same structure in their implementation with regard to `invoke` and `if` expressions, which makes them narrowing handlers.

Figure 4.4 illustrates a narrowing handler refining the contract shown in Figure 4.3. The handler implements an additional requirement for the figure editor example that “some figures are fixed and thus they may not be changed or moved”. To implement the constraint the field `fixed` is added to the class `Fig`, line 23. For fixed figures the value of this field is 1 and 0 otherwise. The class `Point` is the same as in Figure 1.1. To implement the constraint the handler `check` skips invoking the base code whenever the figure is fixed (checked by accessing the field `fixed`).

```

12 class Enforce{
13   Fig check(thunk Fig rest, Fig fe){
14     if(fe.fixed == 0)
15       invoke(rest)
16     else
17       refining establishes fe==old(fe){
18         fe
19       }
20   }
21   when Changed do check;
22 }

```

```

23 class Fig { int fixed; }

```

Figure 4.4 Narrowing handler

For the handler `check` to refine the contract in the event type `Changed`, its implementation must structurally match the contract. The true block of the `if` expression on line 14–15 refines the true block

of the **if** on lines 5–6 as they textually match. The false block of the **if** on line 16–19 refines the false block of the **if** on lines 7–8 because lines 17–19 claim to refine the specification on line 8. This claim is discharged by runtime assertions.

4.3 Direct Interference: Replacement

A replacement handler omits the execution of the original event body and runs the handler body instead. In Ptolemy this can be achieved by omitting the **invoke** expression in the handler.

```

1 Fig event Moved{
2   Point p;
3   int d;
4   requires p != null && d > 0
5   assumes{
6     preserves p != null && p.y == old(p.y)
7   }
8   ensures p != null
9 }

```

No invoke

Figure 4.5 Specifying replacement with a translucent contract

Figure 4.5 shows the contract in event type `Moved` specifying replacement handlers by not evaluating any **invoke** expression in the contract, line 6. Notice that (non) existence of an **invoke** expression in the contract *requires* the handler implementation to (not) evaluate the **invoke** in its body.

```

10 class Scale{
11   int s;
12   Fig scaleit(thunk Fig rest, Point p, int d){
13     refining preserves p!=null && p.y==old(p.y){
14       p.x += s*d; p
15     }
16   }
17   when Moved do scaleit;
18 }

19 class Point extends Fig{
20   int x, int y;
21   Fig moveX(int d){
22     announce Moved(this, d){
23       this.x += d; this
24     }
25   }
26 }

```

Figure 4.6 Replacement handler

Figure 4.6 shows a replacement handler refining the contract in Figure 4.5. The example uses several standard sugars such as `+=` and `>`. In this example, the method `moveX` causes a point to move along the x-axis by amount `d`. The handler `scaleIt` implements the requirement that the “amount of movement should be scaled by a scaling factor `s`, defined in class `Scale`”.

If an contract has no **invoke** expression, none of the refining handlers are allowed to have an **invoke** in their implementation. Otherwise the structural similarity criterion of the refinement is violated. The handler `scaleIt` refines `Moved`’s contract because its body on lines 13–15 matches the specification on line 6.

4.4 Direct Interference: Combination

A Combination handler, typically useful for fault tolerance, can functionalities, can evaluate **invoke** expression any number of times. (In AspectJ, this would be equivalent to one or more calls to **proceed** in an around advice, guarded by some condition or in a loop.) Figure 4.7 illustrates a combination contract and a handler. The translucent contract in the event type specification on lines 5–11 allows an **invoke** expression to be evaluated zero or more number of times. This is achieved by guarding the **invoke** expression by **while**. Based on the contract specially looking at the while loop surrounding `invoke`, the base code developer can conclude that handler methods for event `ClChange` may run the original event body multiple times. The developer, however, is not aware of the concrete details of handlers, thus those details remain hidden.

A combination handler is illustrated in Figure 4.7 lines 15–34. In this example, colors are added to the figures elements by adding a field `color` to the class `Fig` and by providing a method `setColor` for picking the color of the figure, lines 35–43. The class `Color` which provides a method `nextCol` to get the next available color is not shown.

To implement the requirement that “each figure should have a unique color”, event type `ClChange` is declared as an abstraction of events representing colors changes. The method `setColor` changes colors so it announces the event `ClChange` on lines 39–41. The body of the announce expression contains the code to obtain the next color on line 40. The handler `Unique` on lines 15–34 implements this requirement by storing already-used colors in a hash table (`colors`). The field `colFix` is added

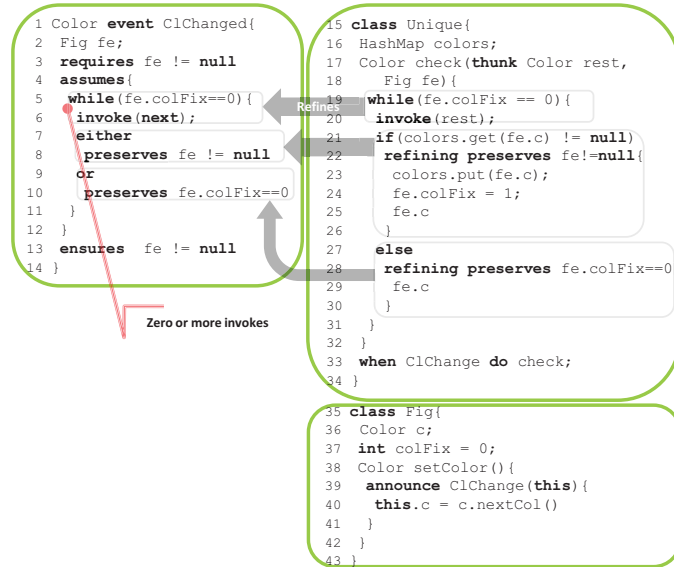


Figure 4.7 Combination contract and handler

to class `Fig` to show that a unique color has been chosen and fixed for the figure. When the handler method `check` is run it checks `colFix` to see if a color has been chosen yet or not. If not then it invokes the event body generating the next candidate color. If the color is already used, checked by looking it up in the hash table, event body is invoked again to generate the next candidate color. Otherwise, the current color is inserted into the hash table and `colFix` is set to 1, lines 21–26.

The specification for `ClChange` on lines 4–12 says that a combination handler will be run when this event is announced. The specification makes use of the choice feature, on line 7–10. To correctly refine the specification, based on the refinement rules in Figure 3.2, a handler can either have a refining `if` expression at the corresponding place in its body or it can have an unconditional expression refining the either-block or the or-branch in the specification. Refinement between specification and implementation blocks is illustrated in the figure.

4.5 More Expressive Control Flow Properties

Rinard *et al.*'s control flow properties are only concerned about calls to `invoke`. Their proposed technique decides which class of interference and category of control effects each isolated advice belongs to [24]. However, it can not be used to analyze the possibility of two or more control flow paths

each of which being, e.g. an augmentation, if each path maintains a different invariant. Figure 4.8 illustrates such a scenario with an example adapted from [12].

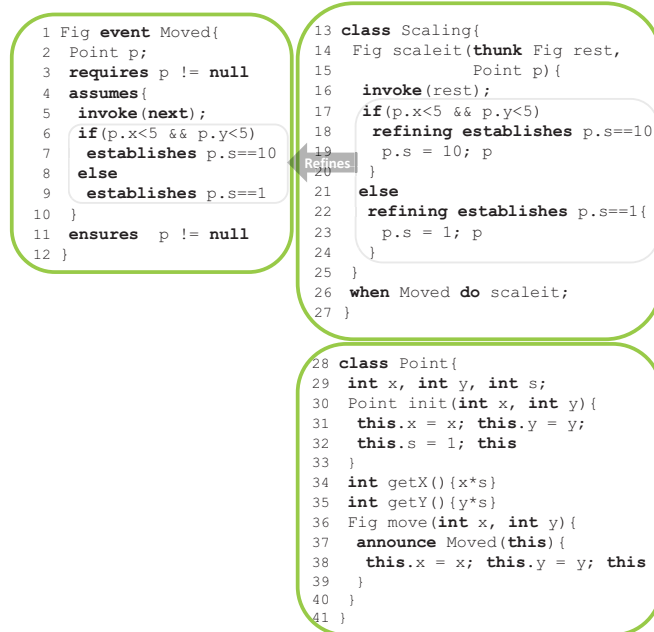


Figure 4.8 Expressive control flow properties beyond [24]

In this example the requirement is “a point should be visibly distinguished from the origin” [12]. If the point is close enough to the origin, its coordinates will be scaled up by a scaling factor s added to `Point` on line 29, initially set to 1, line 32. The scaling factor s has only two values: 1 and 10. The requirement is implemented in the handler method `scaleit` which runs whenever event `Moved` is announced and sets up the scaling factor to 10 if the point is close enough to the origin (vicinity condition). The vicinity condition is true if the point’s x and y coordinates are both less than 5. The class `Fig` is the same as in Figure 1.1.

The assertions to be validated here are as follows: (i) all of the handlers are after-augmentation ones, (ii) the scaling factor s is either 1 or 10, and (iii) s is set to 10 *if and only if* the vicinity condition holds. Rinard *et al.*’s proposal could only be used to verify (i) and a behavioral contract could specify (ii) but none of them could specify (iii). However translucent contracts can. On lines 6–9 there is a specification that conveys to the developer of the class `Point` that a conforming handler method will satisfy all three of the above-mentioned assertions.

In summary, we showed how translucent contracts enable specification and automatic verification, via structural refinement, of control flow interference between a subject and its observers. Translucent contracts are expressive enough to specify and enforce Rinard *et al.*'s [24] control interference even ones which could not be specified by previous works on the design by contract for aspects.

CHAPTER 5. Applicability to Other AO Interfaces

We now discuss the applicability of our technique to other approaches for AO interfaces. As discussed previously, there are several competing and often complementary proposals for AO interfaces. For example, Kiczales and Mezini’s aspect-aware interfaces (AAI) [13], Sullivan *et al.*’s crosscutting interfaces (XPIs) [28], Aldrich’s Open Modules [1], and Steimann *et al.*’s join point types [27]. We have tried out several of these ideas and our approach works beautifully. Since Steimann *et al.*’s join point types [27] and Hoffman and Eugster’s explicit join points (EJP) are similar in spirit to Ptolemy, which we have already discussed in previous chapters, we do not present the straightforward adaptation of our ideas to their work here. Rather we focus on the AspectJ implementation of the XPI approach [28], Kiczales and Mezini’s AAIs [13], and Aldrich’s Open Modules [1] that are substantially distinct from event types [19, Fig. 10].

5.1 Translucid Contracts for XPIs and AAIs

Sullivan *et al.* [28] proposed a methodology, that they call crosscut programming interface (XPI) for aspect-oriented design based on design rules.

```

1 aspect Changed {
2   pointcut jp(Fig fe):
3     call(void Fig+.set*(..)) && target(fe);
4   requires fe != null
5   assumes{
6     if(fe.fixed == 0)
7       proceed(fe);
8   else
9     establishes fe == old(fe);
10  }
11 ensures fe != null
12 }

```

Figure 5.1 Applying translucid contract to an XPI

The key idea is to establish a design rule interface which serves to decouple the base design and the aspect design. These design rules govern exposure of execution phenomena as join points, how they are exposed through the join point model of the given language, and constraints on behavior across join points (e.g. provides and requires conditions [28]).

XPIs prescribe rules for join point exposure, but do not provide a compliance mechanism. Sullivan *et al.* have shown that at least some design rules can be enforced automatically using AspectJ's features[28]. Current proposals for XPIs, however, all use behavioral contracts [28].

As shown previously, use of behavioral contracts, limits the expressiveness of the assertions which could be made using XPI. Behavioral contracts cannot reveal control flow details, which might be needed for reasoning about interference from control effects in cases such as those discussed above.

In this section, we show that translucent contracts can also be applied to enable expressive assertions about aspect-oriented programs that use the XPI approach. We also discuss changes in the refinement rules that are needed to verify such programs. To illustrate, consider the narrowing example from Section 4.2 shown in Figure 5.1 and Figure 5.2, where the constraint on movement of figures is implemented as an XPI and an aspect. Figure 5.1 shows the XPI `Changed` along with the translucent contract on lines 4–11. An XPI typically also contains a description of scope, which we omit here. In the context of XPIs, the language for expressing translucent contract is slightly adapted to use **proceed** instead of **invoke** on line 7. Other than that, our syntax works right out-of-the-box.

```

13 aspect Enforce {
14   Fig around(Fig fe): Changed.jp(fe){
15     if(fe.fixed == 0)
16       proceed(fe);
17     else
18       refining establishes fe==old(fe){
19         return fe;
20       }
21   }
22 }

23 class Fig { int fixed; }

```

Figure 5.2 Narrowing advice for XPI

Unlike translucent contracts for event types in Ptolemy, where the contract is thought of as attached to the type, in the XPI, contracts are thought of as attached to the pointcut declaration, e.g. the contract

on lines 4–11 is attached to the pointcut on lines 2–3. The variables that can be named in the contract are those exposed by the pointcut. For example, the contract can only use `fe`.

Our proposal for verifying refinement also needs only minor changes. Figure 5.2 shows a refining advice for the translucent contract of Figure 5.1. Unlike Ptolemy, where the event types of interest are specified in the binding declarations, in Sullivan *et al.*'s version of XPIs, aspects reuse the pointcut declarations from the XPI in the advice declaration (lines 14). Our refinement rules could be added here in the AO type system. So for an advice declaration to be well-formed, its pointcut declaration must be well-formed, the advice body must be well-formed, and the advice body must refine the translucent contract of the pointcut declaration. This strategy works for basic pointcuts, for compound pointcuts constructed using rules such as `(pcd1 && pcd2` or `pcd1 || pcd2)`, where both `pcd1` and `pcd2` are reused from different XPIs and thus may have independent contracts more complex refinement rules will be needed, which we have not explored in this thesis.

```

1 Point extends Fig {
2   int x, int y;
3   Fig setX(int x): Update -
4   after returning Update.jp(Fig fe)
5     requires fe != null
6     assumes{
7       if(fe.fixed == 0)
8         proceed(fe);
9     } else
10      establishes fe == old(fe);
11  }
12  ensures fe != null
13  /* body of setX */
14 }

```

Figure 5.3 Applying translucent contract to an AAI

Join point interfaces like XPIs could be computed from the implementation rather than being explicitly specified, given whole-program information. Kiczales and Mezini [13] follow this approach to extract aspect-aware interfaces (AAI). A detailed discussion of the trade-offs of such interfaces is the subject of previous work[28]. However, an important property of AAIs is that advised join points contain the details of the advice. An example based on the narrowing example of Section 4.2 is shown in Figure 5.3. The extracted AAI for the method `setX` is shown on lines 3-4. An adaptation of this extraction to include translucent contracts will be to carry over the contract from the pointcut to the join point shadow as shown on lines 5–12.

Syntax and refinement rules similar to XPIs are applicable here. Like AAI annotations that provide developers of `Point` with information about potentially advising aspects, added contract would provide developers of `Point` with richer abstraction over the aspect's behavior. Similar ideas can also be applied to aspect-oriented development environments such as AJDT, which provide AAI-like information at join point shadows in an AO program.

5.2 Translucid Contracts for Open Modules

Aldrich's proposal on Open Modules [1] is closely related to Ptolemy's quantified, typed events [19]. Open Modules allows a class developer to explicitly expose pointcuts for behavioral modifications by aspects, which is similar to signaling events using the **announce** expressions of Ptolemy. The implementations of these pointcuts remain hidden from the aspects. As a result, the impact of base code changes on the aspect is reduced. However, quantification in Ptolemy is more expressive compared to Open Modules. In open modules, each explicitly declared pointcut has to be enumerated by the aspect for advising. On the other hand, Ptolemy's quantified, typed events significantly simplify quantification. Instead of manually enumerating the join points of interest, one can use the name of the event type for implicit non-syntactic selection of join points. This affects applicability of translucid contracts to Open Modules.

```

1 module FigModule {
2   class Fig;
3   expose to Enforce: call(void Fig+.set*(..));
4   requires fe != null
5   assumes{
6     if(fe.fixed == 0)
7       proceed(fe);
8     else
9       establishes fe == old(fe);
10  }
11  ensures fe != null
12 }

```

Figure 5.4 Applying translucid contract to Open Modules

To show the applicability of translucid contracts to Open Modules, we revisit the narrowing example from Section 4.2. Figure 5.4 and Figure 5.5 show the implementation of the same scenario using Open Modules. In implementing the example, we use the syntax from the work of Ongkingco *et al.*

```

13 aspect Enforce {
14 Fig around(Fig fe): target(fe) &&
15 call(void Fig+.set*(..));
16 if(fe.fixed == 0)
17   proceed(fe);
18 else
19   refining establishes fe==old(fe){
20     return fe;
21   }
22 }
23 }
24 class Fig { int fixed; }

```

Figure 5.5 Narrowing handler for Open Module

[16] to retain similarity with other examples. In the listing constraints on the movement of figure is encapsulated in the module (aspect) `Enforce` in Figure 5.5. Open module `FigModule` in Figure 5.4 exposes a pointcut of `class Fig` on line 2–3, marked by the keyword **expose to**. The exposed pointcut is advisable only by the aspect `Enforce`. The translucent contract on lines 4–11 states the behavior of interaction between specified aspect `Enforce` as shown in Figure 5.5 and the exposed pointcut through **expose to** construct. The adaptations in the syntax of contracts are the same as in the case of the XPIs discussed in Section 5.1.

Like contracts in XPIs, contracts in Open Modules are attached to a pointcut declaration, e.g. the contract on lines 4–11 is attached to the exposed pointcut defined on lines 2–3. Variables that can be named in the contract are those exposed by the pointcut, e.g., the contract can only use the variable `fe`.

The rules proposed for verifying refinement need to be modified slightly as well. In Ptolemy, event type of interest is specified in the binding declaration whereas in AspectJ’s version of Open Modules, aspects could not reuse pointcuts exposed by an Open Module and need to enumerate the pointcut in the advice declaration again (lines 14–15). Our refinement rules could be added here in an AO type system. Well-formedness of basic and compound pointcuts follow the same rules laid out in Section 5.1.

This example illustrates how our approach might be used as a specification and verification technique for Open Modules. The only challenge that we saw in this process was to match an aspect’s pointcut definition with the open module’s pointcut definition to import its contract for checking refinement. Like translucent contracts for Ptolemy, in the case of Open Modules specification serves as a more expressive documentation of the interface between aspects and classes.

CHAPTER 6. Related Ideas

There is a rich and extensive body of ideas that are related to ours. Here, we discuss those that are closely related under three categories: contracts for aspects, proposals for modular reasoning, and verification approaches based on grey box specification.

6.1 Contracts for Aspects

This work is closest in the spirit to the work on crosscutting programming interfaces (XPIs) [28]. XPIs also allow contracts to be written as part of the interfaces as **provides** and **requires** clauses. Similar to translucent contracts, the **provides** clause establishes a contract on the code that announces events, whereas the **requires** clauses specifies obligations of the code that handles events. However, the contracts specified by these works are mostly informal behavioral contracts and thus are not easily checked automatically. Furthermore, these works do not describe a verification technique and contracts could be bypassed.

Skotiniotis and Lorenz [26] propose contracts for both objects and aspects in their tool *Cona*. *Cona*'s contracts are black box, and thus do not reveal any information about control flow effects.

Similarly, *Pipa* is a behavioral specification language for AspectJ [31]. *Pipa* supports specification inheritance and specification crosscutting. It relies on textual copying of specifications for specification inheritance and syntactical weaving of specification for specification crosscutting. AspectJ program annotated with JML-like *Pipa*'s specifications could be transformed into JML and Java code. JML-based verification tools could enforce specified behavioral constraints. All of these ideas use black box contracts and thus may not be used to reason about control effects of advice.

6.2 Modular Reasoning

There is a large body of work on modular reasoning about AO programs on language designs [1, 6, 9], design methods [13, 28], and verification techniques [10, 14]. Our work complements ideas in the first and the second categories and can use ideas in the third category for improved expressiveness. Compared to work on reasoning about implicit invocation [3, 8], our approach based on structural refinement is significantly lightweight. Furthermore, it accounts for quantification that these ideas do not account for.

Oliveira *et al.* [17] introduce a non-oblivious core language with explicit advice points and explicit advice composition requiring effects modeled as monads to be part of the component interfaces. Their statically typed model could enforce control and data flow interference properties. Their work shares commonalities with ours in terms of explicit interfaces having more expressive contracts to state and enforce the behavior of interactions. However, it is difficult to adapt their ideas built upon their non-AO core language, to II, AO, and Ptolemy as they do not support quantification.

Hoffman and Eugster’s explicit join points [9] and Steimann *et al.*’s join point types [27] share similar spirit with Rajan and Leavan’s event types [19]. Although Steimann *et al.* proposed informal behavioral specification, their work has no explicit notion of formally expressed and enforced contracts, or stating interaction behavior, nor do any of these other approaches.

The work of Khatchadourian *et al.* [11] is closely related in that it addresses both specification and modular verification of AO programs. They use a rely-guarantee approach to specification and verification. Black box behavioral specifications are attached to PCDs in pointcut interfaces, in a way similar to our work. The **assumes** part of a translucent contract plays a role similar to the rely conditions in their specifications, since it specifies the possible state transformations that advice may implement. Structural refinement in our approach plays a role similar to the guarantee part of their specification, since it also limits what the advice (or handler) can do. The main difference is that they use “join point traces” to reason about control effects, which adds an extra burden on the specifier and verifier compared to our grey box approach, which allows more traditional reasoning about control effects in terms of the underlying programming language’s control flow. Their approach is based on black box behavioral specification.

6.3 Grey Box Specification and Verification

This work builds upon previous research on grey box specification and verification [5]. Among others, Barnett and Schulte have used grey box specifications written in AsmL [4] for verifying contracts for .NET, Wasserman and Blum [30] also use a restricted form of grey box specifications for verification, Tyler and Soundarajan [29] and most recently Shaner *et al.* [25] have used grey box specifications for verification of methods that make mandatory calls to other dynamically-dispatched methods. Rajan *et al.* have used grey box specification to enable expressive assertions about web-services [23]. Compared to these ideas, our work is the first to consider grey box specification as a mechanism to enable modular reasoning about code that announces events and handles events, which is a common idiom of AO and II languages.

CHAPTER 7. Soundness of Reasoning

To reason about a method's body (e) containing **announce** and **invoke** expressions, we use the translation algorithm shown in Figure 3.4 to generate a simulating specification expression (se) (see Chapter 3). We claim that the method body expression e is a Hoare logic-based refinement of generated simulating specification expression se [25]. In other words, if starting with a precondition state ϕ_p the specification expression se implies the postcondition state ϕ_q , then starting with the same precondition state ϕ_p and by running e , we will reach the postcondition state ϕ_q . This condition is formalized in the definition below.

Definition 1 (*Hoare Logic Refinement*) *A specification expression se is said to be Hoare-logic-refined by expression e , expressed as $se \lesssim e$, if and only if for all predicates over program states ϕ_p and ϕ_q ,*

$$\phi_p\{se\}\phi_q \Rightarrow \phi_p\{e\}\phi_q.$$

To prove our claim, we rely on Shaner *et al.*'s work on reasoning about object-oriented programs that contain specification expressions [25]. This work proves that an object-oriented program expression e_{oo} is a Hoare-logic refinement of an object-oriented specification expression se_{oo} , if e_{oo} 's structure matches se_{oo} 's structure and for every specification expression $spec$ in se_{oo} there is a corresponding **refining** expression in e_{oo} that claims, and is verified to, refine $spec$ according to Hoare logic. We incorporate their result as the lemma below.

Lemma 1 (*Shaner-Leavens-Naumann Soundness*) *Let se_{oo} and e_{oo} be specification and program expressions and let $se_{oo} \sqsubseteq e_{oo}$, as defined in Figure 3.2, then for all predicates over program states ϕ_p and ϕ_q ,*

$$\phi_p\{se_{oo}\}\phi_q \Rightarrow \phi_p\{e_{oo}\}\phi_q.$$

But Shaner *et al.* only prove their results for object-oriented expressions (meaning the expressions in their paper [25]). To apply these results to reasoning about Ptolemy programs, we must reduce

both Ptolemy-specific specification expressions and program expressions to object-oriented expressions (from [25]). Below we give some sub-results along those lines.

Lemma 2 shows that the *translation* algorithm (Figure 3.4) produces object-oriented (OO) specification expressions whereas lemma 3 shows that the *substitution* algorithm, of Figure 7.3, produces OO program expressions. The translation algorithm replaces invoke expressions by event type contract, whereas the substitution algorithm replaces invoke expressions by the body of the next applicable handler in the chain of handlers to simulate event announcement.

Lemma 2 (*Translation Produces Object-Oriented Specification Expressions*) *Let se_{pt} be an expression which may contain Ptolemy-specific expressions and let se_{oo} be the result of the application of applying the translation algorithm shown in Figure 3.4 to se_{pt} , i.e. $se_{oo} = Tr(se_{pt}, \mathbf{skip}, \perp)$. Then se_{oo} is an object-oriented specification expression.*

The proof of this lemma is trivial and is done by cases on the translation algorithm.

In previous work, Rajan and Leavens [20] have developed a semantics of Ptolemy programs where Ptolemy-specific expressions are natively supported. For the purpose of soundness proof here, consider an alternative version of operational semantics as shown in Figure 7.1. In Ptolemy's alternative semantics, execution of announce and register expressions result in the execution of program expressions which are the result of application of substitution algorithm to announce and register expressions respectively. The substitution algorithm replaces invoke expressions with the body of the next handler in the chain of handlers, thus rules for invoke expressions originally found in Ptolemy's operational semantics [19] are not needed anymore. Ptolemy's original semantics uses a list of active objects A to keep track of registered observer objects, in alternative semantics presented here a constant memory location loc_A in the store, which points to an object which stores list of active objects. The (REFINING) rule, along with (EVALBODY) and (EVALPOST), make sure that a refining expression truly refines the implementation it hides and claims to refine. Aside from the changes described here, the rest of the Ptolemy's operational semantics remains the same as originally proposed in [19].

The alternative operational semantics along with lemma 3 pave the way to conclude that substitution algorithm applied to announce expressions produces a program expression which simulates the

Evaluation relation: $\hookrightarrow: \Gamma \rightarrow \Gamma$

$$\begin{array}{c}
\text{(REGISTER)} \\
\frac{e' = \text{Subst}(\mathbf{register}(e), \mathbf{skip}, \perp, \mathbf{null})}{\langle \mathbf{register}(e), J, S \rangle \hookrightarrow \langle e', J, S \rangle} \\
\\
\text{(ANNOUNCE)} \\
\frac{e' = \text{Subst}(\mathbf{announce } p(\bar{e})\{e\}, \mathbf{skip}, \perp, \mathbf{null})}{\langle \mathbf{announce } p(\bar{e})\{e\}, J, S \rangle \hookrightarrow \langle e', J, S \rangle} \\
\\
\text{(REFINING)} \\
\frac{n \neq 0}{\langle \mathbb{E}[\mathbf{refining } \mathbf{requires } n \mathbf{ ensures } e\{e'\}], J, S \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{evalbody } e' e], J, S \rangle} \\
\\
\text{(EVALBODY)} \\
\frac{\rho = \text{envOf}(\nu) \quad t = \Pi(v) \quad \rho' = \Pi \uplus \{\mathbf{result} : v\} \quad \Pi' = \text{tenvOf}(\nu) \quad \Pi' = \Pi \uplus \{\mathbf{result} : \mathbf{var } t\} \quad \nu' = \mathbf{lexframe } \rho' \Pi'}{\langle \mathbb{E}[\mathbf{evalbody } v e], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[\mathbf{under } \mathbf{evalpost } ve], \nu' + \nu + J, S \rangle} \\
\\
\text{(EVALPOST)} \\
\frac{n \neq 0}{\langle \mathbb{E}[\mathbf{evalpost } v n], J, S \rangle \hookrightarrow \langle \mathbb{E}[v], J, S \rangle} \\
\\
\text{(UNDER)} \\
\frac{\langle \mathbb{E}[\mathbf{under } v], \nu + J, S \rangle \hookrightarrow \langle \mathbb{E}[v], J, S \rangle}{}
\end{array}$$

Figure 7.1 Alternative operational semantics of Ptolemy[20]

behavior of event announcement. Lemma 3 shows that Ptolemy-specific program expressions are reduced to object-oriented expressions using the substitution algorithm.

Lemma 3 (*Substitution Produces Object-Oriented Program Expressions*) *Let loc_A be a constant memory location in store which points to the list of active objects. Let e_{pt} be a program expression which may contain Ptolemy-specific expressions and let e_{oo} be the result of the application of the substitution algorithm shown in Figure 7.3 to e_{pt} , i.e. $e_{oo} = \text{Subst}(e_{pt}, \mathbf{skip}, \perp, \mathbf{null})$. Then expression e_{oo} is an object-oriented program expression.*

Proof of this lemma is again trivial and could easily be carried out by case analysis like lemma 2.

7.1 Substitution Algorithm

The substitution and translation algorithms are similar on one hand, in the sense that they both replace announce and invoke expressions, on the other hand, they are different as substitution algorithm produces a program expression by replacing announce and invoke expressions, whereas translation algorithm results in a specification expression. The translation algorithm replaces announce and invoke expressions with either the event type's contract or the event body, depending on the existence of applicable handlers. The substitution algorithm replaces those expressions with either body of the next handler or event body, again based on the existence of applicable handlers. $\text{Subst}(e, b_e, p, loc_h)$ is the application of substitution algorithm to program expression e , with event p announced and event body

b_e . Instead of list of active objects A in Ptolemy's original semantics, the substitution algorithm uses a constant memory location loc_A . Location loc_A points to an object of class `ActiveList`, which is responsible for tracking the list of receiver objects for applicable handlers.

Most cases of substitution algorithm $Subst$ are straightforward; like those of the translation algorithm, they recursively apply $Subst$ to each subexpression and compose the results. Figure 7.3 shows how to do that. For Ptolemy-specific expressions, the rule for **refining** $spec\{e\}$ basically applies the substitution algorithm to the subexpression e . The rule for **register**(e), first applies the substitution algorithm to the subexpression e and then adds it to the list of the applicable handlers. The most interesting cases are those for the invoke and announce expressions. In the substitution of these expressions specially for invoke expression, the assumption is that the contract for event type p is of the form **requires** sp_p **assumes** $\{se_p\}$ **ensures** sp'_p . Consequently in the substitution of announce expression the contract for event p' will be like **requires** $sp_{p'}$ **assumes** $\{se_{p'}\}$ **ensures** $sp'_{p'}$.

In both cases conditional if expressions are produced as the body of a refining expression. The refining expression claims to refine the black box behavioral specification $spec$ of the event type p . The refinement of the specification expression by the body of a refining expression is taken care of by run time assertion checking, as discussed in Section 3.4.5.

$Subst(\mathbf{invoke}(e), b_e, p, loc_h)$ produces a conditional if expression which checks for the number of applicable handlers. In its true branch, the conditional expression, contains a sequence of two expressions: substitution of parameter expression e and substitution of the event body b_e , with the assumption that there are no more applicable handlers. Likewise, the false branch of the conditional contains a sequence of two expressions: result of the substitution of parameter expression e and result of the substitution of the body of the next applicable handler. The assumption of this branch is the existence of more applicable handlers. Compare this to the translation of invoke expression in Section 3.4.3.

In case of an announce expression **announce** $p'(\bar{e})\{e\}$, the result of substitution is again a conditional if expression checking for the number of applicable handlers. The true branch of the conditional contains a sequence of two expressions: substitution of parameter expressions \bar{e} and substitution of the event body e . The assumption in this branch is that there are no more applicable handlers. The false branch of the conditional contains a sequence of two expressions as well: result of the substitution of

parameter expression \bar{e} and result of the substitution of the body of the next applicable handler. Readers are encouraged to compare this to the translation of announce expression in Section 3.4.3.

Figure 7.4 shows auxiliary functions used in the substitution algorithm. Function $suc(loc_h, p)$ returns the body of the next handler of event p using the location loc_h which points to the list of applicable handlers for event p . The function gets the location of the first handler of event p by calling method $getFirst()$ and performs a standard β -reduction on the handler method's body. α -renaming takes care of name clashes, if any. Auxiliary function $findHandler(c, p, CT)$ returns the handler for event p in class c where CT is a list of program declarations. Function $eventsOf(CT, loc)$ returns a list of events that object loc observes.

```

1 class ActiveList {
2   Hashtable hash;
3   LinkedList handlers(Event p){
4     LinkedList hList = null;
5     hList = (LinkedList)hash.get(p); hList
6   }
7   void add(Object o, Event p){
8     LinkedList hList = null;
9     hList = (LinkedList)hash.get(p);
10    if(hList != null)
11      hList.add(o)
12    else{
13      hList = new LinkedList();
14      hList.add(o);
15      hash.put(p, hList)
16    }
17  }
18  void add(Object o, LinkedList evs){
19    Event p = null;
20    int size = evs.size();
21    for(int i=0 ;i<size; i++){
22      p = (Event)evs.remove(i);
23      add(o,p)
24    }
25  }
26 }
27 class HashTable {...}
28 class LinkedList {...}
29 class Event {...}

```

Figure 7.2 Classes to simulate list of active objets

To implement the substitution algorithm we assume the existence of some pre-defined classes like `ActiveList` as shown in Figure 7.2. `ActiveList` keeps track of the list of active objects per event type. Handlers of each specific event are stored in a `LinkedList`. Constant location loc_A points to an object of type `ActiveList`. Method `add(Object o, LinkedList evs)` in `ActiveList`

adds object `o` as the observer for all events in the list `evs`. Classes `Hashtable` and `LinkedList` are the same as classes `Hashtable` and `LinkedList` in Java. Class `LinkedList` has an extra method `tail` which returns the tail of the list.

If b_e : event body, p and p' : event types \bar{t} \overline{var} context variables for p and \bar{t}' \overline{var}' context variables for p' translucent contract for p is: requires sp_p assumes $\{se_p\}$ ensures sp'_p translucent contract for p' is: requires $sp_{p'}$ assumes $\{se_{p'}\}$ ensures $sp'_{p'}$ loc_A : Constant location for list of active objects loc_h : Location for the list of handlers of event p Then $Subst(e, b_e, p, loc_h) =$		
Cases of e	Result	Side Conditions
$n, \text{new } c(), \text{var } \text{null}$	e	
$e.f$	$e'.f$	if $e' = Subst(e, b_e, p, loc_h)$
$e.m(\bar{e})$	$e'.m(\bar{e}')$	if $e' = Subst(e, b_e, p, loc_h)$, $\bar{e}' = Subst(\bar{e}, b_e, p, loc_h)$
$e_0.f = e_1$	$e'_0.f = e'_1$	if $e'_0 = Subst(e_0, b_e, p, loc_h)$, $e'_1 = Subst(e_1, b_e, p, loc_h)$
if (e){ e_0 } else { e_1 }	if (e'){ e'_0 } else { e'_1 }	if $e' = Subst(e, b_e, p, loc_h)$, $e'_0 = Subst(e_0, b_e, p, loc_h)$, $e'_1 = Subst(e_1, b_e, p, loc_h)$
while (e) { e_0 }	while (e') { e'_0 }	if $e' = Subst(e, b_e, p, loc_h)$, $e'_0 = Subst(e_0, b_e, p, loc_h)$
cast c e	cast c e'	if $e' = Subst(e, b_e, p, loc_h)$
$e_0; e_1$	$e'_0; e'_1$	if $e'_0 = Subst(e_0, b_e, p, loc_h)$, $e'_1 = Subst(e_1, b_e, p, loc_h)$
$t \text{ var} = e_0; e_1$	$t \text{ var} = e'_0; e'_1$	if $e'_0 = Subst(e_0, b_e, p, loc_h)$, $e'_1 = Subst(e_1, b_e, p, loc_h)$
refining spec { e }	refining spec { e' }	if $e' = Subst(e, b_e, p, loc_h)$
register (e)	$loc_A.add(e', evs)$	if $e' = Subst(e, b_e, p, loc_h)$, $evs = eventsOf(CT, e')$
invoke (e)	refining spec { if ($k == 0$){ e' ; b_e } else { e' ; e'' } }	if $loc_h = loc_A.handlers(p)$, $k = loc_h.size()$, $spec = \text{requires } sp_p \text{ ensures } sp'_p$, $e' = Subst(e, b_e, p, loc_h)$, $loc_h^{tail} = loc_h.tail()$, $e'' = Subst(suc(loc_h, p), b_e, p, loc_h^{tail})$
announce p' (\bar{e}) { e }	refining spec { if ($k' == 0$){ \bar{e}' ; e' } else { $\bar{t}' \overline{var}' = \bar{e}'$; e'' } }	if $loc_{h'} = loc_A.handlers(p')$, $k' = loc_{h'}.size()$, $spec = \text{requires } sp_{p'} \text{ ensures } sp'_{p'}$, $\bar{e}' = Subst(\bar{e}, b_e, p, loc_h)$, $e' = Subst(e, b_e, p', loc_h)$, $loc_{h'}^{tail} = loc_{h'}.tail()$, $e'' = Subst(suc(loc_{h'}, p), e', p', loc_{h'}^{tail})$

Figure 7.3 Substitution algorithm

7.2 Proof of Soundness

To prove the soundness of our reasoning approach, we have proved the translation algorithm sound, i.e., that the specification expression produced by translation algorithm used for reasoning is refined by

$$suc(loc_h, p) = \begin{cases} e_\beta & \begin{array}{l} \mathbf{if} \text{ } loc_h \neq \mathbf{null}, \\ \mathbf{where} \text{ } loc = loc_h.getFirst(), \\ [c.F] = S(loc), \\ th(\mathbf{thunk} \ t' \ var_0, \bar{t} \ \overline{var})\{e_h\} = \\ \quad findHandler(c, p, CT), \\ e_\beta = e_h[\mathbf{this}/loc] \end{array} \\ \mathbf{null} & \mathbf{if} \text{ } loc_h == \mathbf{null} \end{cases}$$

Figure 7.4 Auxiliary functions of substitution algorithm

the program expression produced by substitution algorithm. Theorem 1 formalizes this.

To reason about a method which may announce an event, translation algorithm is applied to the method body, e_{pt} , which may include Ptolemy-specific expressions and the result specification expression se_{oo} is used to reason about the method. Lemma 2 assures se_{oo} is an OO specification expression and therefore can be used for reasoning purposes based on Shaner *et al.*'s approach [25] as stated by lemma 1. This is possible only, if there is a guarantee that se_{oo} is specifying the runtime behavior of the method. The substitution algorithm along with the alternative operational semantics given in Figure 7.1 simulates the original Ptolemy's operational semantics for event announcement. Lemma 3 makes sure that the result of the application of substitution algorithm to e_{pt} is an OO program expression, e_{oo} . Finally theorem 1 guarantees that se_{oo} is stating the behavior of e_{oo} , i.e $se_{oo} \lesssim e_{oo}$, definition 1.

Theorem 1 (Refinement Theorem) *Let program expression e be the body of a method m and $se' = Tr(e, \mathbf{skip}, \perp)$ be the translation of e . Let $e' = Subst(e, \mathbf{skip}, \perp, \mathbf{null})$ be the substitution of e . Then: $se' \lesssim e'$.*

Proof: The proof is by induction on the cases of expression e . For each case we prove $se' \sqsubseteq e'$ as defined in Figure 3.2 and then conclude $se' \lesssim e'$ based on lemma 1 and definition 1. Proof given here is based on the cases of e where e is a non-specification expression. Thus specification expressions **next**, **old** (se), **either** $\{se\}$ **or** $\{se\}$, **requires** sp **ensures** sp are not considered in the proof.

- $e \in \{n, \mathbf{var}, \mathbf{null}, \mathbf{new} \ c()\}$, this is vacuously true because $se' = e$ and $e' = e$ and any expression is refined by itself, i.e, $e \sqsubseteq e$. Therefore $se' \sqsubseteq e'$ which in turn implies $se' \lesssim e'$ based on lemma 1 and definition 1.

- $e = e.m(\bar{e})$, where $se' = Tr(e.m(\bar{e}), \mathbf{skip}, \perp)$ and $e' = Subst(e.m(\bar{e}), \mathbf{skip}, \perp, \mathbf{null})$.
Based on the induction hypothesis a subexpression in se' is refined by its corresponding subexpression in e' . And based on the definition of the translation and substitution algorithms it is easy to see that se' and e' are structurally similar. Therefore $se' \sqsubseteq e'$.
- For $e \in \{e.f, e.f e = e, \mathbf{if}(e)\{e\} \mathbf{else}\{e\}, \mathbf{cast} c e, e; e, \mathbf{while}(e)\{e\}, t \mathbf{var} = e; e\}$, the proof is similar to the proof for method call case of $e = e.m(\bar{e})$.
- $e = \mathbf{refining} spec\{e\}$, where $se' = spec$ and $e' = \mathbf{refining} spec\{Subst(e, \mathbf{skip}, \perp, \mathbf{null})\}$. Refining expression e' is refining specification expression $spec$ which is the same as se' .
- $e = \mathbf{register}(e)$, based on the induction hypothesis a subexpression in se' is refined by its corresponding subexpression in e' . As it can be seen the substitution of register expression is manipulating the list of active objects through loc_A . An Unrolling strategy in the specification expression generated by translation algorithm takes care of different number of handlers.
- $e = \mathbf{invoke}(e)$, again induction hypothesis assures a subexpression in se' is refined by its corresponding subexpression in e' . Also recall that each handler method refines its event type specification which means refinement of se_p by the body of the next handler $suc(loc_h, p)$. Structural similarity of se' and e' could easily be seen in Figure 7.5. Translation and substitution of invoke and announce expressions is shown in this figure. Refinement rules in Figure 3.2 assure either-or block on translation side for invoke expression in Figure 7.5 is refined by if-else block on substitution side.
- $e = \mathbf{announce} p'(\bar{se})\{se\}$. Based on the induction hypothesis a subexpression in se' is refined by its counterpart subexpression in e' . Structural similarity of se' and e' could easily be seen in Figure 7.5. ■

$Tr(\mathbf{invoke}(se), b_e, p)$	$Subst(\mathbf{invoke}(e), b_e, p, \mathbf{null})$
<pre> refining spec{ either {se'; b_e} or {se'; se''} } where : se' = Tr(se, b_e, p) and se'' = Tr(se_p, b_e, p) spec = requires sp_p ensures sp'_p </pre>	<pre> refining spec{ if(k == 0){e'; b_e} else{e'; e''} } where : loc_h = loc_A.handlers(p) and k = loc_h.size() and spec = requires sp_p ensures sp'_p e' = Subst(e, b_e, p, loc_h) and loc_h^tail = loc_h.tail() and e'' = Subst(suc(loc_h, p), b_e, p, loc_h^tail) </pre>
$Tr(\mathbf{announce } p'(\overline{se})\{se\}, b_e, p)$	$Subst(\mathbf{announce } p'(\overline{e})\{e\}, b_e, p, \mathbf{null})$
<pre> refining spec{ either {se'; se'} or {l' var' = se'; se''} } where : se' = Tr(se, b_e, p) and se'' = Tr(se_p', se', p') spec = requires sp_p' ensures sp'_p' </pre>	<pre> refining spec{ if(k' == 0){e'; e'} else{l' var' = e'; e''} } where : loc_h' = loc_A.handlers(p') and k' = loc_h'.size() and spec = requires sp_p' ensures sp'_p' e' = Subst(e, b_e, p, loc_h) and e'' = Subst(suc(loc_h', p), e', p', loc_h'^tail) </pre>

Figure 7.5 Structural similarity of translation and substitution of announce and invoke expressions

Proving theorem 1 means our proposed reasoning approach is sound. In other words statically computed translation of a Ptolemy expression containing announce and invoke expressions, is an object-oriented specification expression which could be used for reasoning purposes without being dependent on runtime configuration of the system, i.e. number of the handlers and their order of execution.

CHAPTER 8. Conclusion and Future Work

We showed how to modularly specify and verify Ptolemy programs that use dynamically announced events and handlers, which is similar to AspectJ's pointcuts and dynamic advice.

First, Ptolemy [19] provides a notion of event type declarations. Event announcement names an event type, and so code announcing an event can use the translucent contracts given in the event type declaration. Similarly, handlers are statically bound to event types in *binding* declarations, and this allows binding verification to also modularly refer to the event type's translucent contract. As the interface between event announcements and handlers, event type declarations are thus a good place to write translucent contracts. We also demonstrated the applicability of our techniques to other type of AO interfaces [1, 9, 13, 27, 28]. Second, Ptolemy's explicit announcement solves the problem of frequent join point shadows, since one only has to deal with handlers where events are explicitly announced. Finally, and most importantly, using grey box specifications as part of our translucent contracts, and using structural refinement in verification solves the problem of reasoning about control effects of handlers. In essence, the grey box specification exposes all the interesting control effects of handlers and structural refinement ensures that correct handler implementations are limited to the specified control effects. We argued that black box behavioral contracts are insufficient for reasoning about such control flow effects, but showed how our translucent specifications were adequate to specify a wide variety of such control effects. We have added translucent contracts to a Ptolemy compiler that verifies handler refinement and inserts runtime assertion checking code [18].

Adding translucent contracts to other AO compilers, integrating our ideas with the rich specification features of JML, and working out larger examples to find out more of the practical use cases of translucent contracts are some directions for future work. Another direction is to use translucent contracts to reason about *data effects* of subject-observer interaction patterns.

BIBLIOGRAPHY

- [1] J. Aldrich. Open modules: Modular reasoning about advice. In *ECOOP '05*.
- [2] M. Bagherzadeh, H. Rajan, and G. T. Leavens. Translucid contracts for aspect-oriented interfaces. In *FOAL '10*.
- [3] L. Baresi, C. Ghezzi, and L. Mottola. On accurate automatic verification of publish-subscribe architectures. In *ICSE '07*.
- [4] M. Barnett and W. Schulte. Runtime verification of .NET contracts. *Journal of Systems and Software*, 65(3), 2003.
- [5] M. Büchi and W. Weck. The greybox approach: When blackbox specifications hide too much. Technical Report 297, Turku Center for Computer Science, August 1999.
- [6] C. Clifton and G. T. Leavens. MiniMAO₁: Investigating the semantics of proceed. *SCP '06*, 63(3).
- [7] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for Java. In *PLDI '02*.
- [8] D. Garlan, S. Jha, D. Notkin, and J. Dingel. Reasoning about implicit invocation. In *FSE '98*.
- [9] K. J. Hoffman and P. Eugster. Bridging Java and AspectJ through explicit join points. In *PPPJ '07*.
- [10] S. Katz. Diagnosis of harmful aspects using regression verification. In *FOAL '04*.
- [11] R. Khatchadourian, J. Dovland, and N. Soundarajan. Enforcing behavioral constraints in evolving aspect-oriented programs. In *FOAL '08*.

- [12] R. Khatchadourian and N. Soundarajan. Rely-guarantee approach to reasoning about ao programs. In *SPLAT '07*.
- [13] G. Kiczales and M. Mezini. Aspect-oriented programming and modular reasoning. In *ICSE '05*, pages 49–58.
- [14] S. Krishnamurthi, K. Fisler, and M. Greenberg. Verifying aspect advice modularly. In *FSE '04*.
- [15] J. M. Morris. A theoretical basis for stepwise refinement and the programming calculus. *Sci. Com. Program.*, 9(3), 1987.
- [16] N. Ongkingco *et al.*. Adding Open Modules to AspectJ. In *AOSD '06*.
- [17] B. Oliveira, T. Schrijvers, and W. R. Cook. Effective advice: Disciplined advice with explicit effects. In *AOSD '10*.
- [18] Ptolemy with Translucid Contracts. <http://www.cs.iastate.edu/~ptolemy/contract/>.
- [19] H. Rajan and G. T. Leavens. Ptolemy: A language with quantified, typed events. In *ECOOP '08*.
- [20] H. Rajan and G. T. Leavens. Quantified, typed events for improved separation of concerns. Technical Report 07-14, Iowa State University, Department of Computer Science, July 2007.
- [21] H. Rajan and K. J. Sullivan. Classpects: unifying aspect- and object-oriented language design. In *ICSE '05*.
- [22] H. Rajan and K. J. Sullivan. Unifying aspect- and object-oriented design. *TOSEM '08*.
- [23] H. Rajan, J. Tao, S. M. Shaner, and G. T. Leavens. Tisa: A language design and modular verification technique for temporal policies in web services. In *ESOP '09*.
- [24] M. Rinard, A. Salcianu, and S. Bugrara. A classification system and analysis for aspect-oriented programs. In *FSE'04*.
- [25] S. M. Shaner, G. T. Leavens, and D. A. Naumann. Modular verification of higher-order methods with mandatory calls specified by model programs. In *OOPSLA '07*.

- [26] T. Skotiniotis and D. H. Lorenz. Cona: Aspects for contracts and contracts for aspects. In *OOPSLA '04*.
- [27] F. Steimann, T. Pawlitzki, S. Apel, and C. Kastner. Types and modularity for implicit invocation with implicit announcement. *TOSEM '10*, 20(1).
- [28] K. J. Sullivan, W. G. Griswold, H. Rajan, Y. Song, Y. Cai, M. Shonle, and N. Tewari. Modular aspect-oriented design with XPIs. *TOSEM '09*, 20(2).
- [29] B. Tyler and N. Soundarajan. Black-box testing of grey-box behavior. In *FATES '03*, 1–14.
- [30] H. Wasserman and M. Blum. Software reliability via run-time result-checking. *J. ACM*, 44(6):826–849, 1997.
- [31] J. Zhao and M. Rinard. Pipa: A behavioral interface specification language for AspectJ. In *FASE '03*.