# User Queries for Specification Refinement Treating Shared Aspect Join Points

Emilia Katz, Shmuel Katz
*Computer Science Department*
*Technion – Israel Institute of Technology*
*Haifa, Israel*
*Email: {emika, katz}@cs.technion.ac.il*

***Abstract*—We present an interactive semi-automatic procedure to help users refine their requirements formally and precisely, using knowledge the user possesses but does not notice as relevant and has difficulty formalizing. Questions in natural language are presented to the user, and augmentations to specifications, written in Linear Temporal Logic, are automatically created according to the answers. We apply our approach to a case study on specifying the desired aspect behavior in a delicate case when multiple aspects can share a join-point, i.e., be applied at the same state of base program computation. The questions used in the case study are derived from an in-depth analysis of semantics and mutual influence of aspects at a shared join-point. Aspects sharing a join-point might, but do not have to, semantically interfere. Our analysis and specification refinement enables programmers to distinguish between potential and actual interference among aspects at shared join-points, when aspects are modeled as state transition diagrams, and specifications are given as LTL assumptions and guarantees. The refined aspect specification, obtained from the procedure we describe, enables modular verification and interference detection among aspects even in the presence of shared join-points.**

***Keywords*-formal specification; semantics; aspect interference; shared join-points**

## I. Introduction

Writing a specification in a precise and formal way is not an easy task, as it is hard for the user to think of all the relevant details. After the specification is almost completed, there still might be some delicate special cases not treated by it, not from the lack of information, but from the fact that the user might not notice that this information is indeed relevant, and also might not immediately know how to formalize it. The goal of our paper is to connect informal knowledge of the user with the input expected by formal verification tools, especially model checkers, verifying Temporal Logic specifications. We present a semi-automatic interactive procedure for specification refinement, that will help the user refine the specification in such cases.

In the procedure we propose, the user is asked a series of questions in natural language, regarding the parts of system behavior that are important for the delicate case treated. Based on the answers, statements in Temporal Logic are automatically created, to be added to a previously existing specification. The semi-automatic creation of properties uses parametrized formulae, which are a result of previous analysis of the semantics of the special case in question. In this paper we apply our approach to aspect semantics and mutual influence of aspects at a shared join-point.

Aspect-oriented software development aids in modular extension of object systems, where treatment of concerns that cut across many parts of the system is encapsulated in separate modules called *aspects*. An aspect consists of two parts: the code associated with concern treatment (called *advice*), and a predicate defining when advice should be applied during system executions (called a *pointcut descriptor*, or - more briefly - a *pointcut*). A pointcut can refer both to static and dynamic information about the state of the system, i.e., identify not only the place in the code, or other conditions that can be statically evaluated, but also runtime conditions under which the advice should be invoked (such as values of system variables, or call stack contents). Thus, when a computation arrives at one and the same place in the code several times, sometimes the corresponding state may be matched by the pointcut, and sometimes it may not. The points in the execution that are identified by a pointcut are called *join-points*. The process of combining aspects with a system is called *weaving*, the original system is then referred to as the *base system*, and the result - as the *woven system*. The best-known aspect-oriented language is the AspectJ [1] extension of Java, but there are many aspect-oriented languages and software development techniques (see, e.g., [2]).

When multiple aspects are woven into the same base system, the question of interactions among aspects arises. One possible situation, widely recognized as potentially problematic, is the case when different aspects woven into the same base system happen to have common join-points. This possibility gives rise to several key questions and problems, from understanding how the potentially applicable advice pieces should be woven at such a common join-point (as they cannot be applied all at the same time), to conflict detection and resolution, since application of one advice might interfere with the computation or even applicability of another.

Application of our approach to this delicate situation results in an easily automatizable interactive procedure. It helps users specify their intentions for aspect behavior in a specific system more precisely, and check whether there

is actual interference with respect to this specification. We model aspects and systems as state transition diagrams, with specifications given as linear temporal logic (LTL) assumptions and guarantees. We analyze the possible influences of multiple aspects on each other at a join-point, and classify them according to the way they affect the specification of the aspect. Based on this analysis, natural language questions to the user are manually constructed. According to the answers to these questions, the LTL specifications of the aspects are automatically augmented. Then the modular verification procedure from [3], [4] and modular automatic interference checks from [5] can be used to verify the correctness of each aspect alone and to detect cases of interference at shared join-points, or establish that there is no such interference, using the augmented specifications. This case study has been presented in a workshop ( [6]), with emphasis on the aspect verification details rather than the user-guided specification refinement procedure.

The paper is organized as follows: First we discuss related work in Section II. Our analysis of aspect semantics and types of mutual influence at shared join-points appears in Section III. In Section IV we show the questions the answers to which can be automatically processed to add appropriate predicates to a temporal logic specification. Section V gives examples of applying the questions. We conclude in Section VI.

## II. Related Work

Ways to detect shared join-points are described in [7], [8]. Several works study shared join-points as a source of possible conflicts, some (e.g., [7]) even see common join-points as the main source of interference among aspects. A language independent technique [9] makes it possible to check whether an undesired order of aspect application at a shared join-point is possible, where the list of undesired orders has to be explicitly provided by the user. It is implemented in the "Secret" tool for Compose* [10], [11]. However, presenting the list of undesired orders requires a thorough analysis of the system, and also might not reflect all the intended behaviors, as at different states different orders of application might be possible. In [12] another tool for checking potential interference at common join-points is described, applicable for the Compose* language. It checks all the possible orders of aspect applications at a common join-point, and declares a conflict if different orders result in different resulting states. This method is fully automatic, but may lead to many false positives (some of which are described later in this paper). An additional tool for aspect interference detection, performing dataflow checks to find out whether one aspect affects variables used in another, is presented in [13]. It can also be used to check interactions at shared join-points, though its scope is broader. Interactions found by this tool are, like the previous tools, only potentially harmful.

Weaving techniques for conflict resolution at shared join-points appear in [7], [14], [15]. In [7], a first analysis of types of mutual influence of aspects applied at a shared join-point appears. This analysis is extended in our paper, though used for a different purpose.

The intended semantics of weaving several aspects at a common join-point in a pre-defined order is addressed in papers on the semantics of aspects, such as [16], [17], [18].

However, as described in [19], not all the conflicts at shared join-points can be resolved by a clever weaving. Thus it is important for the user to be able to detect the conflicts and differentiate between real problems and false alarms.

## III. Semantics of aspect behavior at a common join point

In AspectJ, the most commonly used aspect language, aspects at a common join-point are applied one after another, and each time before performing an advice the pointcut condition is re-checked. As a result of such a semantics, when a base system arrives at a join-point matched by an aspect A, it is not necessarily the case that the advice of A is immediately executed. Other aspects in the system might also match this join-point, and thus some other advices might be executed before the advice of A, changing the state of the system in which A will be applied. Moreover, A's advice might not be executed at all, in case one of the previously executed aspects left the system in a state which is no longer a join-point of A.

Thus the execution of the woven system from the moment it arrives at a join-point matched by some of its aspects is determined not only by the set of matching aspects, but also by the order of their application at this point. If this order of application is not explicitly prescribed by the user, the non-determinism of aspect application may result in different states. However, the fact that different orders of advice application lead to different resulting states does not necessarily mean that the aspects semantically interfere. Consider the following example, from [12]: Several aspects are defined for systems in which messages of type String are sent between objects. Two of these aspects are Logging and Encryption. Both aspects are applied at the same join-points - when a message is sent in the system - and different orders of their application will result in different states of the system. If Logging is executed before Encryption, the logged message will be the original one, otherwise it will be the encrypted message produced by the Encryption aspect. In [12] such a situation is considered interference between the two aspects, but in fact the decision on whether it is interference should depend on the aspects' specifications. In our example, the goal of the Encryption aspect is to encrypt every message before it is sent to the server. Consider the following possible specifications of the Logging aspect, described more formally in Section V:

1) The log should record all the messages sent as they were originally written by the user, so that the user can view the contents of the messages.
2) The log should record all the messages as they were actually sent to the server in order to compare the sent messages to the received ones (as received) and verify that no messages got lost or garbled.
3) The goal of the Logging aspect is to measure the network activity of the system. Thus, though the contents of the messages are written to the log, they are of no importance to the user, and what matters is only, e.g., the times of the messages sent and the number of lines in the log.
4) The logging records all the attempts to send a message, even if they are aborted for whatever reason. It logs each message as it was attempted to be sent.

All the cases above can happen in our example system, and different orders of application of the two aspects at their common join-point lead to different resulting states, but not in all the cases above do the aspects interfere. The requirements from the Encrypting aspect are never violated by Logging, no matter in what order they are executed, but in variants (1) and (4), the goal of Logging will not be reached if the Encrypting aspect is applied first, but the opposite order is permissible. In variant (2), only applying Encrypting after Logging will cause a problem. Moreover, in variant (3) applying the aspects in any order will not violate the requirements from Logging or from Encrypting, thus there is no interference. As will be shown later, an Authorization aspect can also be applied, further complicating the situation.

The above example shows the need to analyze possible semantic effects of sharing a join-point more deeply. We consider the AspectJ operational semantics. Under it, the weaving is performed by a three-step strategy: First all the places in the code of the base program that are matched by the static part of some aspect's pointcut, are identified. Such places are called *shadow join-points*. Note that a shadow join-point is usually defined by a place in the code of the program, but sometimes can contain additional information. Second, after shadow join-point identification, at each such join-point the weaving order of the potentially applicable aspects is defined (an aspect is considered potentially applicable if the static part of its pointcut matches the current shadow join-point). The weaving order can potentially be determined dynamically, upon arrival of the computation at a join-point. Finally, when a computation arrives at a join-point, each of the potentially applicable aspects, one by one and in the previously defined order, is checked for full applicability and immediately executed if indeed applicable (i.e., if both static and dynamic parts of the pointcut are matched by the current state). All the rest of the paper refers to this semantics, and if a different semantics is chosen, different reasoning might be needed. This operational semantics shows the need to reason about the part of computation between the first moment it arrives at some shadow join-point and the moment it leaves this join-point, which includes all the aspect applications performed at the join-point. We need some new terminology to make this reasoning easier. First of all, we need a name for the period of interest:

*Definition 1:* A sequence of states $s_1, \ldots, s_k$ in a computation of the woven system is called a *pointcut occurrence* of aspect A if $s_1$ is the state when a join-point of A is first reached (that is, $s_1$ is matched by the full pointcut descriptor of A, and the previous state is not), and $s_k$ is the state when the computation is about to leave the corresponding shadow join-point, after application of all the appropriate aspect advices according to the current weaving policy (that is, $s_k$ is matched by the static part of the pointcut descriptor of A, and the next state is not).

Two aspects share a join-point if they have at least one overlapping pointcut occurrence. Note that overlapping pointcut occurrences do not have to coincide, as it might be a case that an execution arrives at a state $s$ matched by the pointcut descriptor of aspect B, and by the static part of pointcut descriptor of A and B, but not matched by the dynamic part of A's pointcut descriptor, and only the execution of aspect B at $s$ will result in a state in which both static and dynamic parts of A's pointcut hold. In such a case the pointcut occurrence of A will be contained in the pointcut occurrence of B, but not vice versa. For example, consider the case of aspects A and B applied to a grades managing system. Let aspect B be responsible for giving bonuses and factors to grades, and let aspect A be in charge of enforcing a required grades format, by rounding non-integer grades and replacing all the grades above 100 by 100. Both aspects are applied before the publication of the grades, so the static parts of their pointcuts are the same. However, aspect A should be applied only if the grade to be published is not an integer or exceeds 100. Clearly, a computation might arrive at a place before grade publishing with an integer grade below 100, thus matched by the dynamic part of B's pointcut only (and not of A's), but as a result of B's modifications, a non-integer grade or a grade above 100 is obtained, bringing the computation to a state that is matched by A's pointcut as well.

Previously, two kinds of join-points have been examined: shadow join-points and actual join-points. A shadow join-point of aspect A, as mentioned above, is a place in the code of the base program that is matched by the static part of A's pointcut. An actual join-point of A is a state in a computation of the system at which the advice of A is actually applied. However, for the purpose of our analysis, a third type, *arrival join-points*, is needed:

*Definition 2:* A state $s$ in a computation of the woven system is called an *arrival join-point* of aspect A if $s$ is the first state of a pointcut occurrence of A - the state when a

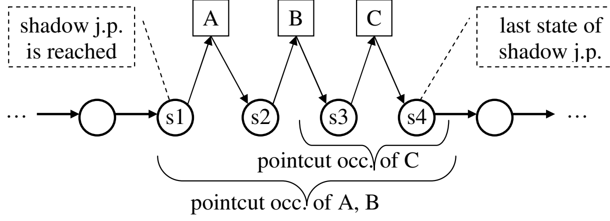

Figure 1. Pointcut occurrence example.

join-point is first reached in that occurrence.

Note that an arrival join-point differs from a shadow join-point because it is matched also by the dynamic part of A's pointcut descriptor. It also differs from an actual join-point because other aspects can intervene and even prevent A from reaching an actual joinpoint.

Figure 1 presents an example for the above definitions. The state $s_1$ is a shadow join-point of aspects A, B and C. It is also matched by the dynamic parts of A's and B's pointcuts, thus both for A and for B the sequence $s_1, \ldots, s_4$ is a pointcut occurrence, and $s_1$ is an arrival join-point. However, their actual join-points differ: A is applied at $s_1$, and B - at $s_2$. For aspect C, the sequence $s_3, s_4$ is a pointcut occurrence, meaning that the dynamic part of the pointcut of C becomes true only after B is executed. $s_3$ is also the state where C is executed, thus being both the arrival and the actual join-point of C.

The pointcut of A identifies states either before or after some events of interest, but the definitions above are applicable for both cases. And if A is an `around` advice (i.e., advice that augments the event of interest, and sometimes even replaces the functionality of the base system), it either has a `proceed` statement, and then can be viewed as a combination of two advice pieces - one before, and one after the corresponding event, or A has no `proceed`, and then can be viewed as a `before` advice, one of the effects of which is a change in the program counter of the base system. If indeed A's advice changes the program counter of the base system, the end of its execution is also the end of the current pointcut occurrence - both according to our intuition and to the definitions above.

Now let a state $s$ be a join-point matched by aspect A, that appears inside pointcut occurrence $\pi$. We distinguish between four possible cases of other aspects' behavior that can influence the result of weaving A into a system:

1) Aspect B executed before A in $\pi$ changes a value of some variable used by A as an input to its computations.
2) Aspect C executed after A in $\pi$ changes a value of some variable updated by a computation of A.
3) Aspect D executed before A in $\pi$ brings the system to a state $s'$ which is not a join-point of A any more.
4) Aspect E executed after A in $\pi$ invalidates the condition on which the join-point predicate depended, thus removing a join-point of A after A has already been executed at it.

The following analysis of those cases enables us to determine whether the above influences actually cause an interference: Let the specification of A be given by an assumption-guarantee pair $(P_A, R_A)$, where $R_A$ is the guarantee of A that must hold in any woven system containing A, provided the system into which A has been woven satisfied the assumption of A, $P_A$. Note that A can be woven into a system that does not satisfy $P_A$, but then $R_A$ is not guaranteed to hold in the resulting system. We denote by $V_{in}(A)$ a set of variables A uses as input to its computations, and by $V_{out}(A)$ - a set of variables in which A stores the result of its computations.

**1. Change Before (CB).** In case an aspect B executed before A at $s$ changes a value of a $v \in V_{in}(A)$, the result of A's calculations might differ from the one we would get if the value of $v$ has not been changed from the moment the computation arrived at $s$ till the moment the advice of A was executed. If the guarantee of A is formulated in terms of a specific connection between the value of $v$ when we arrive at a join-point and the value of $v$ after the computation of A is finished, $R_A$ will be violated. (This can happen, for instance, in variant (1) of the Logging and Encrypting example above: we anticipate that the message string written to the log is the one created by the user and readable by the user, but if Encrypting is executed before Logging, what we actually get in the log is the encrypted message, because the contents of the message was changed by the Encrypting aspect.) Note that if the requirement for correctness of A's calculations binds the values at the end of A's execution only to the values at the beginning of execution of A (as in variants (2) and (3) of the Logging and Encrypting example, with Encryption before Logging), it will not be violated in this case (for variant (2) only the final message contents are important, and for variant (3) message contents are not important at all).

**2. Change After (CA).** In case some aspect C executed after A at $s$ changes a value of a $v \in V_{out}(A)$, the guarantee of A will be violated if it required preservation of the result of A's computation till some future point in the execution where the value of $v$ is used. (As in variant (2) of the example, when Logging occurs before Encryption). Otherwise, as in variant (3) of the example with Logging before Encryption, the guarantee of A will not be influenced. If, indeed, a requirement for preservation of the value of $v$ till some state *use_v* is part of A's guarantee, then part of A's assumption should be that in the base system the value of $v$ is not modified from the actual place of application of A's advice till arrival to the *use_v* state.

**3. Invalidation Before (IB).** In this case there is no state in A's pointcut occurrence at which A is executed. Such a situation happens, for example, with Logging and Au-

thorization aspects from Section V when the Authorization aspect is applied before Logging and the authorization of the user fails, thus preventing message sending, and removing the join-point of the Logging aspect. In variant (4) of the Logging specification, this leads to violation of the guarantee of Logging, as a message was prepared for sending and should have been logged, but the Logging aspect never has a chance to be applied, because the authorization failure finishes the pointcut occurrence.

**4. Invalidation After (IA).** In this case A is executed at some point at which it shouldn't have been applied, because when arriving at the point of interest, the weaver "does not know" that the reason for A's application will be removed by one of the aspects coming after A. If the specification of A requires that it is applied only if followed by some event in the future, and this following event is removed by another aspect, then the specification of A is violated. This is the case, for example, in variants (1), (2) and (3) if the Authorization aspect is applied after Logging and the authorization of the user fails. Note that in variant (4), on the other hand, the guarantee of Logging is not violated if Logging precedes Authorization.

## IV. Specification of aspects with possibly shared join-points

### A. Guided Specification Construction

An assume-guarantee specification of aspect A is a pair of LTL formulas, $(P_A, R_A)$, where $P_A$ is the assumption of the aspect about all the base systems into which it can reasonably be woven, and $R_A$ is the guarantee of the aspect, that must hold after A is woven into any system that satisfied the assumption $P_A$. Generally, the basic specification of the aspect is clear, as will be seen in the examples in Section V. The refinement of the basic specification for the case of possibly shared join-points, on the other hand, is not obvious, thus the procedure presented below is useful.

The LTL specifications we use include:

- "$\mathsf{G}\,\varphi$" ("Globally") - meaning that the formula $\varphi$ is true from the current state on.
- "$\mathsf{F}\,\varphi$" ("Finally") - from the current state a state in which $\varphi$ holds can be reached.
- "$\mathsf{O}\,\varphi$" ("Once") - dual to "Finally": a state satisfying $\varphi$ occurred earlier in the computation.
- "$\varphi\ \mathsf{U}\ \psi$" ("Until") - a state in which $\psi$ holds is reached later in the computation, and until then $\varphi$ holds.
- "$\varphi\ \mathsf{W}\ \psi$" ("Weak Until") - almost like "Until", but a state in which $\psi$ holds does not have to be reached. In that case $\varphi$ holds forever from the current state.

We say that a computation satisfies an LTL formula if this formula holds in its first state. From the analysis in Section III the need for the following predicates arises:

- *at(ptc)*: assuming that $ptc$ is the predicate defining A's pointcut, the predicate $at(ptc)$ means that the computation has just arrived at a join-point of A. It is useful for reasoning about what happened in the computation after the moment it arrived at a possibly shared join-point. In fact, this is the predicate marking the arrival join-points of A.
- *after_prev_asp(A)*: this predicate becomes true at the moment the weaver has applied all the aspects that preceded A at the current shadow join-point, according to the algorithm of the current weaver. This predicate is used to refine the definition of A's pointcut because now A should only be applied at states satisfying both $ptc$ and *after_prev_asp(A)*, (which matches the definition of the set of all the actual join-points of A). In addition, the predicate is used in assumptions added to A's specification to express the desired behavior of the system from the moment its computation arrives at a join-point of A till the moment A's advice is actually executed.
- *asp_ret(A)*: this predicate describes the possible return states of the aspect. This is needed for some of the cases below. Typically, the aspect return state has the same control location as the join-point state (the values can change, but not the program counter of the join-point). For the Logging aspect, for example, the base state is actually not changed, and only the log (local to the aspect) is modified. However, it does not have to be so in general. Thus in order to define the *asp_ret(A)* predicate, the user is proposed a default predicate, automatically constructed by the system as described in [4]. The idea of construction is to create a system containing representations of all the possible computations of an aspect from all its possible initial states, without actually applying the aspect to any specific base system (this is done using the MAVEN tool [3] and some built-in functionality of the NuSMV [20] model-checker), and then to build a predicate describing all the states of this system that satisfy the return conditions of the aspect. This default predicate can then be manually modified.

Using the above predicates, all the requirements mentioned in Section III can be expressed, though often not all of the predicates are needed in a given application. Below we present a way to express each of the additional requirements.

The construction of the refined specification can be automatic, but user-guided: questions are presented to the user, and the answers to these questions determine the new requirements. The construction process will thus be as follows:

*Step 1:* Here we will treat the dependency of our aspect, A, on its input variables, in order to find out whether the values of the input variables need to be preserved between the arrival and the actual join-points of A (in order to be able to treat the "change before" case from Section III). The user is asked the following question:

**Q. 1:** Are there any input variables of A for which the advice

of A depends on the value as it is at the arrival join-point and not as it is when the advice of A actually starts its execution?

- If yes, the user should provide a list of variables for which such a dependency exists.
- For each variable $v$ in the list, we add the following $CB$ (for "Change Before") statement to A's assumption:

$$CB(v) = \mathsf{G}[(at(ptc) \land v = V) \rightarrow ((v = V \ \mathsf{W} \ (after_prev_asp(A) \land v = V))]$$

where $V$ is a logical variable keeping the value of $v$ as it was at the arrival to the join-point.

- If there are no variables in the list, nothing is added to the specification of A at this step.

*Step 2:* Here we treat the case when part of the effect of the aspect is modification of some state variables, and this effect should be preserved till some point in the future of the computation. This is important for the "change after" case from Section III. The questions asked here are:

**Q. 2:** Are there any state variables of the system into which A is woven the value of which should be preserved after A's execution is finished? (For example, variables modified by A, or variables that are semantically connected to A's local variables.)

- If yes, the user is asked to fill in a table with two columns: the first column is the name of the variable, $v$, and the second is a state predicate *use_v* describing the state of the woven system until which the value of $v$ should be preserved. For example, for variant 2 of the Logging aspect, that logs messages as they are sent to the server, the message should not change between the moment it has been logged and the moment it is actually sent. Thus the *use_v* predicate will describe the moment of actual sending of the message (see Section V for more details). After the table is filled out, for each variable $v$ with state predicate *use_v* in the table, we add the following $CA$ (for "Change After") statement to the assumption of A:

$$CA(v) = \mathsf{G}[(asp_ret(A) \land v = V) \rightarrow (v = V \ \mathsf{W} \ (use_v \land v = V))]$$

where $V$ is a logical variable keeping the value of $v$ as it was at the end of the execution of A's advice.

- If there are no variables in the list, nothing is added to the specification of A at this step.

*Step 3:* In this step we construct requirements corresponding to the "invalidation before" case in Section III. Before the problem of common join-points in modular verification was considered, there existed an implicit assumption that all the arrival join-points of an aspect are its actual join-points. But when a join-point might be shared, this is not necessarily so, because the join-point can be invalidated; thus an additional explicit assumption of this possibility is needed. The user is asked the following question:

**Q. 3:** Does it have to be that each time an arrival join-point of A is reached, A is eventually executed at it? That is, is it an error if previously executed aspects invalidate the condition for A's application?

- If no, nothing is added to the assumption of A in this step.
- If the answer was "yes", the following $IB$ (for "Invalidation Before") statement is added to the assumption of A:

$$IB \triangleq \mathsf{G}[at(ptc) \rightarrow (ptc \ \mathsf{U} \ (after_prev_asp(A) \land ptc))]$$

*Step 4:* The goal of this step is to enable the verification process to treat the case of "invalidation after" from Section III. We ask the user the following questions:

**Q. 4.1:** Does the reason for a state to be A's join-point lie in the future of the computation? That is, does A's pointcut descriptor refer to any event following the join-point? For example, is the advice of A a "before" advice?

- If no, nothing is added to the assumption of A in this step.
- If the answer was "yes", the next question is asked:

**Q. 4.2:** Is it an error if the advice of A is performed, but the presumably-following event does not follow? (For example, because the future computation was changed by other aspects)

- If the answer is "no", nothing is added to A's assumption in this step.
- If the answer is "yes", the user is required to provide a state predicate, *foll_event*, meaning that the desired following event has just occurred. The user is then prompted to provide some optional restrictions on the values immediately after A's execution, the values at the moment the desired event occurs, and the connections between them (including, for example, value preservation). The restrictions should be given in the form of two predicates: *vals_after_asp* and *vals_at_foll_event*. The default value for both predicates is $true$.
- The following $IA$ (for "Invalidation After") statement is then added to the assumption of A:

$$IA \triangleq \mathsf{G}[(asp_ret(A) \land vals_after_asp) \rightarrow F(foll_event \land vals_at_foll_event)]$$

After the above automatic modifications, the specification constructed both captures the requirements of the user regarding the desired effect of aspect application, and contains sufficient assumptions to make the modular verification results applicable to systems with aspects sharing join-points.

### B. Using the Refined Specification

The refined specification of aspects is used during all the stages of the full specification and verification process for

aspect libraries, where a library of aspects is a collection of reusable aspects grouped together for some common purpose, to be applied together to different systems. Given a library of aspects, two things are important for its usage: correctness of each aspect alone with respect to its assume-guarantee specification, and interference detection among the aspects. The question of possibly shared join-points is already important when the specification of individual aspects is defined. At this stage one of the tools for detection of potential interference at common join-points detection can be run, e.g. [12], and only if a potential interference is detected the specification refinement described in Section IV-A has to be performed for the potentially interfering aspects. (If no tool for potential interference detection is run, all the aspect specifications should undergo the process from Section IV-A, to ensure the soundness of the verification process.)

After all the aspects in the library are specified and augmented as described above, existing tools for modular aspect verification ( [3], [4]) and interference detection ( [5]) can be run.

Verification of aspect correctness with respect to its assume-guarantee specification, using these tools, is a three-stage process. First, a tableau representing all the possible computations of systems satisfying the assumption of the aspect is automatically constructed from the assumption formula (this is done by the `ltl2smv` module of `NuSMV` model-checker). Then the aspect state machine is woven into the tableau, and finally the NuSMV model-checker is run on the resulting woven tableau, to verify the guarantee of the aspect. The tableau contains a representation of all the possible computations of the base system, thus all the possible computations of woven systems containing the aspect are represented in the woven tableau, which means that it is indeed enough to model-check the woven tableau in order to ensure that the guarantee of the aspect will hold in all the computations of any concrete woven system. (The tool for individual aspect verification should undergo some minor technical changes, so that the weaving procedure will ensure that an aspect is never applied more than once at the same join-point)

Interference analysis for the library is based on pairwise interference-freedom checks of the aspects. For each pair of aspects from the library, two properties are verified for each of the two possible weaving orders of the aspects. First, it is verified that when one of the aspects is woven into a system in which both assumptions hold, the assumption of the second aspect will still hold after the weaving (thus enabling the weaving of the second aspect, and guaranteing its correct behavior). Second, it is checked that when one aspect is woven into a system in which the guarantee of the other holds, this guarantee will be preserved after the weaving. Automtization of this verification method uses the above-described tool for modular aspect verification as a subsystem.

Note that the correctness of the individual aspects and of their combinations is verified independently of any concrete base system, thus dividing the whole verification process into two independent parts: whenever an aspect, or a collection of aspects, are to be actually woven into a base system, one part of verification is to check that the base system satisfies the assumptions of all the aspects given, and the other part is to ensure that all the aspects are correct w.r.t. their assume-guarantee specifications, and do not interfere. Such a modularity enables us to check the correctness and interference-freedom of the library of aspects off-line and once and for all, and not each time some aspects are actually woven into a given base system, thus the verification effort is very much reduced. Another advantage of modular verification is that the models verified are smaller, as we never need to actually examine a woven system, and this enhances the model-checking process (and sometimes even makes it possible).

The effect of our specification-refinement procedure on the result of aspect verification and interference detection is twofold. First of all, after the specification of the aspect is refined, it is possible to prove more aspects correct with respect to their assume-guarantee specifications, because the assumption of the aspect might have been strengthened by additional assertions. Second, for the same reason, if two aspects interfere, the interference will in many cases be detected at an earlier stage of the interference-detection procedure, as violation of the assumption of an aspect might be detected, making it unnecessary to proceed to the guarantee-preservation check.

## V. Examples

We illustrate our analysis and specification refinement approach on the collection of aspects described in Section III, that can be a part of a communication-aspects library. The aspects presented are applicable for systems with message-passing, and they are variants of the aspects used as an example in [12].

**Logging aspect (L)** logs the message - sending in the system. As described earlier, there are four variants of the logging aspect in the library:

- $L_1$: Logging all the sent messages as the user originally attempted to send them.
- $L_2$: Logging all the messages that were actually sent to the server (the message is logged as it was sent).
- $L_3$: Logging the frequency of message sending.
- $L_4$: Logging all the attempts to send a message (the message is logged as it was originally attempted to be sent by the user).

Now we need to construct the specifications for the above variants of Logging aspect. The following predicates and variable definitions will be used in the construction:

- *msg_attempt*: a predicate which is true when a message is about to be sent, that is, when message sending is attempted. That is the moment before the message-sending procedure is actually called, and the parameters to the method call are represented by the variables *msg_c* and *msg_t*, containing the two parts of the message to be sent: the contents and the creation time, respectively.
- *msg_send*: a predicate which is true at the moment a message (with contents *msg_c* and creation time *msg_t*) is sent.
- *in_log* ($< str >$) : a predicate that is true if the string "str" appears in the log.

The pointcut of all the variants of the aspect is the moment before the message-sending procedure is called. More formally, $msg_attempt$ is true. The guarantees of the aspects emerge in the usual way from the purpose of each of them, and are written more formally below. If there would be no possibility of sharing join-points, the assumption of all the logging aspect variants would be that if a message is sent, there indeed was an attempt to send this message (i.e., this very same message has been passed as a parameter to the message-sending procedure). More formally,

$$P_L \triangleq \mathsf{G}([msg_send \land msg_c = C \land msg_t = T] \\ \rightarrow \mathsf{O}[msg_attempt \land msg_c = C \land msg_t = T])$$

However, we are aware of the possibility of each of the aspects to share a join-point with other aspects, such as Encryption and Authorization, thus additional assumptions for the aspects are constructed according to the procedure from Section IV-A.

*Specification for $L_1$::* A possible guarantee for $L_1$ is:

$$R_{L1} \triangleq \mathsf{G}([at(msg_attempt) \land msg_c = C \land \\ msg_t = T \land F(msg_send)] \\ \leftrightarrow [\mathsf{F}(in_log(< X, T >))])$$

meaning that messages that appear in the log are all the sent messages, but as they were first attempted to be sent by the user. (Note that the fact that each message is accompanied by creation time information ensures a one-to-one correspondence between messages and lines in the log.)

The answers for the assumption-construction questions for $L_1$ are as follows:

Q.1: "Yes". The aspect depends on the contents and time information of the message as they were at the join-point, thus the values of the variables *msg_c* and *msg_t* should be preserved. Thus, substituting into the template $CB(v)$, the following statements are added to the assumption of $L_1$:

$$CB(c) = \mathsf{G}[(at(msg_attempt) \land msg_c = C) \rightarrow \\ ((msg_c = C) \\ \mathsf{W}\ (after_prev_asp(L1) \land msg_c = C))]$$

and

$$CB(t) = \mathsf{G}[(at(msg_attempt) \land msg_t = T) \rightarrow \\ ((msg_t = T) \\ \mathsf{W}\ (after_prev_asp(L1) \land msg_t = T))]$$

Q.2: "Yes". The time information of the message should be kept intact till the moment the message is actually sent. There is one entry in the table: the variable *msg_t*, matched by the *msg_send* predicate. Thus the addition to the aspect assumption at this stage is:

$$CA(t) = \mathsf{G}[(asp_ret(L_1) \land msg_t = T) \rightarrow \\ ((msg_t = T)\ \mathsf{W}\ (msg_send \land msg_t = T))]$$

Q.3: "No". If the message will not be sent, it should not appear in the log, thus the advice of $L_1$ should not be applied for it. Nothing is added to the assumption of $L_1$ at this stage.

Q.4.1: "Yes". The advice of $L_1$ is a "before" advice.

Q.4.2: "Yes". It is an error if a message that is not sent and will not be sent appears in the log. The desired following event is the event of sending the message (defined by its creation time only, as that is what matters for the purpose of $L_1$). Thus *foll_event* = *msg_send*, *vals_after_asp* = (*msg_t* = *T*) and *vals_at_foll_event* = (*msg_t* = *T*). Substituting into the $IA$ template, we obtain the following addition to $L_1$'s assumption:

$$IA = \mathsf{G}[(asp_ret(L_1) \land msg_t = T) \rightarrow \\ F(msg_send \land msg_t = T)]$$

*Specification for $L_2$::* A possible guarantee for $L_2$ is:

$$R_{L2} \triangleq \mathsf{G}([\mathsf{F}(msg_send\ \land msg_t = T \land msg_c = C)] \leftrightarrow \\ [\mathsf{F}(in_log(< C, T >))])$$

meaning that a message appears in the log if and only if it has been, or will be, sent.

The construction of the assumption for $L_2$ is performed similarly to that of $L_1$, with only two differences: An additional variable, *msg_c*, should be preserved after the aspect finishes its computation (affecting the $CA(v)$ and $IA$ statements), and no values from arrival join-point should be kept (making $CB(v)\ true$). Together we obtain that the addition to the assumption of $L_2$ as a result of the guided specification construction procedure consists of the following statements:

$$CA(c) = \mathsf{G}[(asp_ret(L_2) \land msg_c = C) \rightarrow \\ ((msg_c = C)\ \mathsf{W}\ (msg_send \land msg_c = C))]$$

$$CA(t) = \mathsf{G}[(asp_ret(L_2) \land msg_t = T) \rightarrow \\ ((msg_t = T)\ \mathsf{W}\ (msg_send \land msg_t = T))]$$

and

$$IA = \mathsf{G}[(asp_ret(L_2) \wedge msg_c = C \wedge msg_t = T) \rightarrow F(msg_send \wedge msg_c = C \wedge msg_t = T)]$$

*Specification for $L_3$::* A possible guarantee for $L_3$ is:

$$R_{L3} = \mathsf{G}([\mathsf{F}(msg_send \ \wedge msg_t = T)] \leftrightarrow [\mathsf{F}(in_log(< T >))])$$

meaning that the log contains all the creation-times of the sent messages.

The construction of the assumption for $L_3$ is almost the same as for $L_2$, except for the fact that the value of *msg_ts_c* need not be preserved after the aspect finishes its computation (thus giving the same $CA(t)$ and $IA$ statements as for $L_1$). Thus the addition to the assumption of $L_3$ is

$$CA(t) = \mathsf{G}[(asp_ret(L_3) \wedge msg_t = T) \rightarrow ((msg_t = T) \ \mathsf{W} \ (msg_send \wedge msg_t = T))]$$

and

$$IA = \mathsf{G}[(asp_ret(L_3) \wedge msg_t = T) \rightarrow F(msg_send \wedge msg_t = T)]$$

*Specification for $L_4$::* A possible guarantee for $L_4$ is:

$$R_{L4} \triangleq \mathsf{G}([at(msg_attempt) \wedge msg_c = C \wedge msg_t = T] \leftrightarrow [\mathsf{F}(in_log(< C, T >))])$$

meaning that the log contains exactly the messages attempted to be sent by the user.

The construction of the assumption for $L_4$ is almost the same as for $L_1$, with the following differences only:

- The answers to Question 2.1 and Question 4.1 are negative, as the logged message does not have to be sent, so $CA = true$ and $IA = true$ in this case.
- The answer to Question 3 is positive, as all the message sending attempts should be logged, including those aborted because of authorization failure.

Thus the additions to the assumption of $L_4$ are:

$$CB(c) = \mathsf{G}[(at(msg_attempt) \wedge msg_c = C) \rightarrow ((msg_c = C) \ \mathsf{W} \ (after_prev_asp(L_4) \wedge msg_c = C))]$$

$$CB(t) = \mathsf{G}[(at(msg_attempt) \wedge msg_t = T) \rightarrow ((msg_t = T) \ \mathsf{W} \ (after_prev_asp(L_4) \wedge msg_t = T))]$$

$$IB = \mathsf{G}[at(msg_attempt) \rightarrow (msg_attempt \ \mathsf{U} \ (after_prev_asp(L4) \wedge msg_attempt))]$$

**Encrypting aspect (E)** is responsible for encrypting messages before sending. E should guarantee that each time a message is sent, it is encrypted. In fact, there is more to E: each time a message is received, it is decrypted. But this part is irrelevant to our example, so we'll ignore it here. E's guarantee can be written as:

$$R_E \triangleq \mathsf{G}(msg_send \rightarrow encrypted(msg_c))$$

where the predicate *encrypted(msg_c)* means that the contents of the sent message are encrypted. The assumption of E, constructed by the procedure in Section IV-A, emerges from the fact that the encrypted message value should be preserved till (and if) it is actually sent:

$$P_E = CA(msg_c) = \mathsf{G}[(asp_ret(E) \wedge msg_c = C) \rightarrow (msg_c = C \ \mathsf{W} \ (msg_send \wedge msg_c = C))]$$

**Authorization aspect (A)** ensures that a message is sent to the server only if the current user has the needed permissions to communicate with the server. A's guarantee can be

$$R_A \triangleq \mathsf{G}(msg_send \rightarrow permit_usr_send)$$

where the predicate *permit_usr_send* means that the user has enough permissions to send the message. When constructing the assumption of A, all the answers to the questions asked happen to be negative, thus A does not need to assume anything about the base system, and we can take

$$P_A \triangleq true$$

After the aspects are specified as above, if the usual verification procedure is applied, several cases of interference will be detected, as shown in Figure 2. A cell in the table corresponding to aspects pair $< M; N >$ can have one of the following values:

- "—" means that there is no interference when weaving executing first M and then N at shared join-points in any appropriate system;
- "X" - if the check is irrelevant, here we do not check interference among the aspect and itself.
- Otherwise, there is interference among the two aspects if M is executed before N at shared join-points, and the cause of the interference is written in the cell, according to the classification from Section III: "CB" stands for Change Before, "CA" - for Change After, "IB" - for Invalidation Before, and "IA" - for Invalidation After.

For example, the cell $< E, L_1 >$ is marked by $CB$, meaning that the Encryption aspect, if woven first, invalidates the assumption of the first variant of Logging, and that the violated part of the assumption is related to the "change before" case from Section III: changing parameter values between arrival and actual join-points of $L_1$. And the cell $< L_1, A >$ is marked by $IA$ as the Authorization aspect, when woven after the first variant of Logging, invalidates

| second / first | E | A | L1 | L2 | L3 | L4 |
|---|---|---|---|---|---|---|
| E | X | --- | CB | --- | --- | CB |
| A | --- | X | --- | --- | --- | IB |
| L1 | --- | IA | X | --- | --- | --- |
| L2 | CA | IA | --- | X | --- | --- |
| L3 | --- | IA | --- | --- | X | --- |
| L4 | --- | --- | --- | --- | --- | X |

Figure 2. Interference checks summary.

the part of Logging specification related to the "Invalidation After" case from Section III: removing a join-point of the aspect after its advice has already been applied. Note that we might want to add more than one variant of Logging to a system, and there is no interference among them, provided their log files are different (then instead of one *in_log* predicate we would have four: $in_log_1 \ldots in_log_4$).

## VI. Conclusions

This paper has concentrated on a problematic issue of how to help programmers refine a specification, in delicate situations in which the user possesses some important information, but is unaware of the fact that this information might be useful for a proof. In such a situation, our approach can be used when it is possible to create the relevant statements in a parametric way, based on the user's answers. We demonstrate our approach on the problem of aspect semantics: possible interference that can arise from shared join-points. The questions asked and the results of formal verification should help the user understand the fine points of such interactions, and how they could affect the correctness of their aspect systems.

## References


[1] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, and W. G. Griswold, "An overview of AspectJ," in *Proc. ECOOP 2001*, ser. LNCS 2072, Jun 2001, pp. 327–353, http://aspectj.org.

[2] R. Filman, T. Elrad, S. Clarke, and M. Akşit, Eds., *Aspect-Oriented Software Development*. Boston: Addison-Wesley, 2005.

[3] M. Goldman and S. Katz, "MAVEN: Modular aspect verification," in *Proc. of TACAS 2007*, ser. LNCS, vol. 4424, 2007, pp. 308–322.

[4] E. Katz and S. Katz, "Modular verification of strongly invasive aspects," in *Languages: From Formal to Natural*, ser. LNCS, vol. 5533. Springer, 2009, pp. 128–147.

[5] ——, "Incremental analysis of interference among aspects," in *FOAL '08*. ACM, 2008, pp. 29–38.

[6] ——, "Semantic aspect interactions and possibly shared join points," FOAL'10 workshop, 2010.

[7] I. Nagy, L. Bergmans, and M. Aksit, "Composing aspects at shared join points," in *NODe/GSEM*, 2005, pp. 19–38.

[8] D. H. Nguyen and M. Südholt, "VPA-based aspects: Better support for aop over protocols," in *4th IEEE International Conference on Software Engineering and Formal Methods (SEFM'06)*, 2006, pp. 167–176.

[9] P. Durr, L. Bergmans, and M. Aksit, "Reasoning about semantic conflicts between aspects," in *ADI'06*, 2006, pp. 10–18.

[10] L. Bergmans, "Towards detection of semantic conflicts between crosscutting concerns," in *AAOS Workshop at ECOOP'03*, 2003.

[11] I. Nagy, L. Bergmans, and M. Aksit, "Declarative aspect composition," in *Software Engineering Properties of Languages and Aspect Technologies (SPLAT) Workshop*, 2004.

[12] M. Akşit, A. Rensink, and T. Staijen, "A graph-transformation-based simulation approach for analysing aspect interference on shared join points," in *AOSD*, 2009, pp. 39–50.

[13] N. Weston, F. Taiani, and A. Rashid, "Interaction analysis for fault-tolerance in aspect-oriented programming," in *MeMoT'07*, 2007, pp. 95–102.

[14] R. Douence, P. Fradet, and M. Sudholt, "Composition, reuse, and interaction analysis of stateful aspects," in *Proc. of 3th Intl. Conf. on Aspect-Oriented Software Development (AOSD'04)*. ACM Press, 2004, pp. 141–150.

[15] R. Douence, P. Fradet, and M. Südholt, "A framework for the detection and resolution of aspect interactions," in *GPCE*, 2002, pp. 173–188.

[16] M. Wand, G. Kiczales, and C. Dutchyn, "A semantics for advice and dynamic join points in aspect-oriented programming," *T. on Programming Languages and Systems*, vol. 26, no. 5, pp. 890–910, 2004.

[17] C. Clifton and G. T. Leavens, "MiniMao$_1$: Investigating the semantics of proceed," *Science of Computer Programming*, vol. 63, no. 3, p. 321374, 2006.

[18] S. D. Djoko, R. Douence, and P. Fradet, "Aspects preserving properties," in *PEPM*, 2008, pp. 135–145.

[19] G. Kniesel, "Detection and resolution of weaving interactions," *T. Aspect-Oriented Software Development*, vol. 5, pp. 135–186, 2009.

[20] A. Cimatti, E. Clarke, F. Giunchiglia, and M. Roveri, "NuSMV: a new Symbolic Model Verifier," in *CAV'99*, ser. LNCS 1633. Springer, 1999, pp. 495–499, http://nusmv.itc.it.