Taming aspects with monads and membranes

Ismael Figueroa, Nicolas Tabareau, Éric Tanter

► To cite this version:

Ismael Figueroa, Nicolas Tabareau, Éric Tanter. Taming aspects with monads and membranes. FOAL'13: Foundations of aspect-oriented languages, Mar 2013, Fukuoka, Japan. 10.1145/2451598.2451600 . hal-00808983

HAL Id: hal-00808983

<https://hal.inria.fr/hal-00808983>

Submitted on 8 Apr 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Taming Aspects with Monads and Membranes

Ismael Figueroa

PLEIAD Lab & ASCOLA Group
DCC University of Chile & INRIA
ifigueroa@dcc.uchile.cl

Nicolas Tabareau

ASCOLA Group
INRIA, France
nicolas.tabareau@inria.fr

Éric Tanter*

PLEIAD Lab
DCC University of Chile
etanter@dcc.uchile.cl

Abstract

When a software system is developed using several aspects, special care must be taken to ensure that the resulting behavior is correct. This is known as the *aspect interference problem*, and existing approaches essentially aim to detect whether a system exhibits problematic interferences of aspects.

In this paper we describe how to control aspect interference by construction by relying on the type system. More precisely, we combine a monadic embedding of the pointcut/advice model in Haskell with the notion of membranes for aspect-oriented programming. Aspects must explicitly declare the side effects and the context they can act upon. Allowed patterns of control flow interference are declared at the membrane level and statically enforced. Finally, computational interference between aspects is controlled by the membrane topology. To combine independent and reusable aspects and monadic components into a program specification we use *monad views*, a recent technique for conveniently handling the monadic stack.

Categories and Subject Descriptors: D.3.3 [Programming Languages]: Language Constructs and Features

General Terms: Languages, Design

Keywords aspect-oriented programming, monads, monad views, membranes, non-interference.

1. Introduction

Aspect-oriented programming (AOP) is a programming paradigm to modularize *crosscutting concerns*. We focus on the pointcut/advice model of AOP [6]. Pointcuts are predicates over *join points*, points of interest raised during program execution, to decide application of advice. Advice is an action to be applied when its corresponding pointcut matches. Aspects are modular units that encompass a number of pointcuts and advice.

Although the benefits of aspects are well understood in isolation, there is a substantial body of research devoted to the *aspect interference problem* [2–4, 8, 10, 15, 16] (also known as the *shared join point problem*). Loosely speaking, two aspects interfere if they modify a shared context, or if there is a control flow dependency between them.

Controlling the interactions between aspects is important, since the correctness of the system may be affected by a seemingly harmless change to the system. Existing approaches aim to *detect* situations of aspect interference in AspectJ-like languages. One technique is to translate a program into a state machine which is model-checked for correctness [3]. Another approach relies on a special-purpose control- and data-flow analysis to determine inter-

ference [8].

In this work we provide a mechanism¹ where programmers must use types to specify the allowed interactions between aspects and components – in contrast to post-hoc detection as in other approaches. The main feature of this mechanism is that the type-checker only accepts programs that satisfy the specified interaction constraints between aspects and components.

More specifically, our approach combines (i) a monadic embedding of the pointcut/advice model in Haskell [11]; (ii) a type-based mechanism for reasoning about effects using monads [7]; and (iii) the model of membranes for AOP [14]. The monadic setting requires aspects to explicitly declare in their type the effects that they are able to perform. Allowed patterns of control flow interference are declared at the membrane level and statically enforced. Finally, computational interference between aspects is controlled by the membrane topology. In addition, to enable the development of independent and reusable aspects and monadic components, we use *monad views* [9], a recently developed technique for handling the monadic stack. In summary, the contributions of this paper are:

- Extending the implementation of the Haskell AOP library to use monad views.
- Developing a modular language extension [11] to provide membrane-based AOP semantics to the Haskell AOP library.
- An example of how to use the extended language to control aspect interference relying only on the Haskell type system².

We consider this work as a proof of concept of how to combine monads with membranes to control aspect interference. Future work will focus on a more rigorous formalization of these techniques with relation to aspect interference.

The paper starts with an overview of the Haskell AOP library (Sections 2 and 4), monadic programming (Section 3), and membranes for AOP (Section 5). Then we illustrate our approach with a simple example (Section 6). Familiarity with Haskell and monadic programming is not strictly required to grasp the essence of this paper, but it would surely help to understand all the details. For a brief overview of monadic programming see [9].

2. Aspects à la Haskell

The embedding of aspects in Haskell [11] consists of a small standard Haskell library with full support for the pointcut/advice model. Aspects, pointcuts, and advice are statically typed, and the (unmodified) Haskell type system proves that all pointcut/advice bindings are safe. In addition to being typed, the embedding is also

*É. Tanter is partially funded by FONDECYT project 1110051.
I. Figueroa is funded by a CONICYT-Chile Doctoral Scholarship.

¹ Full implementation and example available at <http://pleiad.cl/haskellaop>

² We use several GHC extensions, which are summarized in the project website.

```

1 data Jp m a b = Jp (a -> m b) a
2 data PC m a b = Monad m => PC (∀ a' b'. m (Jp m a' b' -> m Bool))
3 type Advice m a b = (a -> m b) -> a -> m b
4 data Aspect m a b c d = (LessGen (a -> b) (c -> m d)) =>
5   Aspect (PC m a (m b)) (Advice m c d)
6
7 the Aspect data constructor is hidden, we use a function instead
8 aspect pc adv = Aspect pc adv newAspectHandle
9
10 advice:
11 ensurePos proceed n = proceed (abs n)
12
13 monadic version of sqrt:
14 sqrtM n = return (sqrt n)
15
16 using an aspect:
17 program :: Int -> (AOT I) Int
18 program n = do deploy (aspect (pcCall sqrtM) ensurePos)
19               sqrtM # n

```

Figure 1: Haskell AOP Join Point Model and a small example of open function application (adapted from [11]).

monadic. This means that aspects must be explicit in their types about the side effects they perform, and that we can use standard monadic reasoning techniques.

The main premise to introduce aspects into functional programming is that functions must be subject to aspect weaving. We use the term *open function application* to refer to a function application that emits a join point, opening it to weaving.

2.1 Join Point Model

The pointcut/advice model is comprised of join points, pointcuts, advice, and aspects. In Figure 1 we summarize these elements and present a small example. A join point represents a function application and contains a function of type $a \rightarrow m b$ and the argument of type a . A pointcut is a predicate on join points, which is parameterized by a monad m , denoting the underlying computational effects. The predicate is parametric with respect to join points, and has access to the effects in m . The types a and b are not used in the right-hand side, but are required for type safety [11].

An advice is a function that executes in place of a join point matched by a pointcut. It receives a function, known as *proceed*, and returns a new function of the same type, which not necessarily applies *proceed* internally. An aspect is a first-class value which comprises a pointcut and an advice, and ensures that the binding is safe using the *LessGen* constraint, which requires the type of the pointcut to be *less general* than the type of the advice.

The example shows how the *ensurePos* advice guarantees that all calls to its advised function are made with a positive argument. *pcCall* is a predefined pointcut that matches all open applications of its argument. In addition, we provide a similar pointcut *pcType* based on type signatures, and logical pointcut combinators *pcAnd*, *pcOr* and *pcNot*; user-defined pointcuts are also allowed.

As it can be seen, aspects are deployed using *deploy*, and open function application is done using *#*, that is *sqrtM # n* is woven while *sqrtM n* is not. The AOT monad transformer (see next section) adds the *aspect-weaving* effect, which allows aspect deployment and open function applications.

3. An Overview of Monadic Programming

Monads are a popular mechanism to embed and reason about computational effects in purely functional languages like Haskell. Monad transformers [5] are type constructors that are used as building blocks to modularly create a monad that combines several effects. Each transformer of the composed monad builds a new effect,

and has access to the effects of the underlying monad. A monad constructed with monads transformers is called *monad stack*.

In addition to the AOT transformer³, some standard monad transformers used in this paper are the *state monad transformer* ST , which threads a value with read-write access into a computation; the *reader monad transformer* RT , which provides read access to an environment to enable context-dependent computations; and the *writer monad transformer* WT which provides write-only access to an auxiliary output value that is accumulated over time (e.g. a log). We also use the identity monad I , which has no computational effect and serves as the base of a monadic stack.

3.1 Polymorphism on the Monadic Stack

Monadic components can impose constraints on the monadic stack, using type classes. For instance, if we need state-monad-like access to an integer value we declare the $\text{S}_M \text{ Int } m$ constraint on m . These class constraints can be seen as *families of monads* and imply that the monadic component is polymorphic with respect to the concrete monadic stack, as long as the requirements are satisfied. In this paper we use the S_M , R_M , and W_M type class constraints, which abstract the behavior of the corresponding monad transformers. The ability to deploy aspects and perform open function applications is abstracted in the AOT_M constraint.

3.2 Handling the Monadic Stack

A benefit of using monads is that they present an abstract interface over arbitrary computational effects, which allows for (more) reusable software components. However, when using monad transformers and constraints on the monad stack it is necessary to establish a mechanism to access the effects of each layer. We now briefly describe current mechanisms; for a detailed description see [9].

Explicit Lifting The most primitive solution is to use explicit lifting operations (which are required to define a monad transformer). This however is not convenient for reusability, since each component is coupled (and limited) to a concrete monadic stack.

Implicit Lifting This mechanism relies on type classes and on the polymorphism of the monadic stack to automatically route the operations to the first layer of the monadic stack that satisfies a constraint. For instance, if m is an instance of S_M , then $\text{WT } m$ is also an instance of S_M which implicitly lifts the state-related operations to m . This is the mechanism currently used in Haskell and in [11].

Monad Views Although implicit lifting is flexible and convenient to use, it is impractical when combining multiple instances of the same monad transformer because the implicitly lifted operation is always performed on the first layer that satisfies the constraint. Thus, lower layers with the same effect cannot be accessed and must be explicitly lifted.

Monad views [9] are a novel mechanism to *virtualize* the monad stack. Using views, a programmer can selectively ignore layers of the monadic stack, such that implicit lifting is redirected to the desired remaining layers. In addition, a view can impose restrictions on particular layers – for instance a state layer can be *seen* as a reader layer, providing read-only access to an otherwise writeable component.

In this work we extended our implementation [11] to use monad views to control data-flow interference between aspects, by precisely defining which layers are accessible to aspects (note however that all the described mechanisms can co-exist in a system).

³ Observe that the implementation of the AOT and the MTB (Section 5) monad transformers is identical to that of the standard state monad transformer – we just give them names that reflect their purpose.

4. A First Approach to Interference Control

Initial work regarding aspect interference is done in [11]. In particular, we adapted two techniques from EffectiveAdvice [7]: control flow combinators, and parametricity on the monadic stack.

Control Flow Combinators. The control flow combinators defined in [7] correspond to the categories described by Rinard *et al.* [8]. These combinators can be directly applied in the Haskell AOP library. The categories of control flow behavior are:

- **Augmentation advice** calls `proceed` exactly once and does not modify the arguments or the return value of `proceed`.
- **Narrowing advice** calls `proceed` at most once, and does not modify the arguments or the return value of `proceed`.
- **Replacement advice** completely replaces the join point, and never calls `proceed`.
- **Combination advice** can call `proceed` any number of times with modified arguments or return value.

It could be argued that this classification could be improved. For instance, the advice in the example of Figure 1 can only be implemented as a combination advice, which is too unrestricted. Indeed, in the Haskell AOP library it is possible to define arbitrary *advice combinators* to enforce new kinds of advice [11], for instance advice that is like augmentation advice but can modify the arguments of `proceed`.

Parametricity on the Monadic Stack. We used the technique of parametricity on the monadic stack [7] to define *non-interfering* advice, pointcuts, and programs. That is, we split the monadic stack into two parts: effects available to base computation, and effects available to aspect computation. For instance, a non-interfering advice is defined as

```
type NIAAdvice t a b =  $\forall m$ . Monad (t m)
     $\Rightarrow$  Advice (NIAOT t m) a b
```

where t is the upper part of the split monadic stack, available to aspects; and m is the lower part of the split stack, available to base computation. $NIAOT$ is a variant of the AO_T transformer that supports a split monadic stack. Since an $NIAAdvice$ has to be parametric with respect to m , it effectively cannot access any effects available to base programs.

The technique of parametricity on the monadic stack can be extended to hide arbitrary layers of the monad stack, but components are harder to write and they may become coupled to a particular stack shape. This is why we use monad views [9], which are more flexible and convenient for independent development of reusable components.

5. Membranes for AOP

Execution Levels and Topological Scoping. A particular instance of the aspect interference problem is that of *computational interference*: when the evaluation of an aspect (either pointcut or advice) produces join points that are visible to others. A specific case is self-computational interference—when an aspect advises its own computation—which results in infinite regression.

To address this situation, Tanter proposed the notion of *execution levels* [12]. In execution levels, execution is stratified in a tower in which the flow of control navigates. Given an initial level, join points are always emitted one level above the current level. Aspects are deployed at a specific level and can only affect join points emitted one level below. This way, computation performed by aspects is not observable by aspects that are deployed at the same level, but only by aspects deployed one level above. In addition, the original computation, which corresponds to the last call to `proceed` in the

Figure 2: Topological scoping. (a) tower—corresponds to execution levels; (b) tree; (c) DAG. The lollypop arrow denotes the advising relation between membranes (adapted from [14]).

advice chain, is always executed at the level at which the join point was emitted.

Execution levels can be used to avoid computational interference between different aspects, for instance to reuse off-the-shelf dynamic analyses aspects [13]. The form of *topological scoping* that execution levels provide is very interesting, however, the imposed topology—a tower—is too restrictive.

Flexible Topological Scoping of Aspects. The idea of membranes for AOP was born from the need to generalize execution levels to different topologies. A membrane encapsulates some computation. An operator is provided to spawn computation inside a membrane. The computation that happens inside a membrane produces join points that are only visible to the aspects registered in the *advising membranes*, i.e. membranes that have been explicitly bound to the current membrane. By binding membranes, it is possible to express flexible topologies of computation, because they directly define a directed graph, called the *membrane topology*, with membranes as vertices, and the advising relation as edges. Figure 2 shows three interesting topologies: a tower, which corresponds to execution levels; a tree, and a DAG. We will exploit the tree topology in the example in Section 6.

Because an aspect is registered in a given membrane, its computation (pointcuts, advice) happens inside that membrane, and is only visible to the advising membranes. Infinite regression can therefore happen only if the membrane topology is cyclic.

Beyond Topological Scoping. So far our description of membranes has focused on the topological scoping of aspects, because it is all we need for this work. Topological scoping does not make use of the programmable nature of membranes. Programmability of membranes is exploited in [14] to enforce certain properties on join point emission and reception. For instance, it is possible to encode the restrictions of the control flow combinators described in Section 4, albeit dynamically instead of statically. Here, we are interested in *static*, type-based reasoning about interference.

Bringing Membranes to Haskell AOP. In this work we implemented the MB_T monad transformer as a modular language extension [11] to the Haskell AOP library to provide membrane-based AOP. Any monadic stack with the form $AO_T MB_T m$ provides the membrane-related operations `newMembrane` (creates a new membrane), `register` (registers an aspect in a membrane), `advise` (binds two membranes), and `evalIn` (evaluates a computation in a given membrane).

6. An Illustration of Interference Control

We illustrate how to control aspect interference extending an example from [14]. Consider a web browser application, which processes URLs into HTML documents. To enhance the performance of the browser, an independent cache aspect is used. In addition, to control the memory consumption of both the browser and the

Figure 3: Independent Components: browser, quota advice, and cache advice. Each component imposes its own constraints on the monadic stack to perform its functionality.

Browser. The `browser` (lines 8-13) retrieves an `HtmlDoc` associated to a given `url`, and updates an internal history with the time of access to the `url`⁴. In the composed program, `browserSession` simulates the use of the browser providing several `urls`. This component requires read-write access to a `history` state, as expressed by the $\mathbb{S}_M$ `History m` constraint in its signature.

Cache Advice. This advice (lines 27-36) is a general purpose cache, which uses a dictionary to store already computed results.

⁵ In general `augBef` can pass a value to `augAft`, but we just pass the unit value `()`. That is why the advice type is `Augmentation m a b ()`. The same happens in the augmentation part of `cacheAdv`.

Observe that we need to expose the operation of the cache advice to aspect weaving. To do it we define the `cachePut` function and perform an open application of it in the body of `augAft`. This requires us to add the $\mathbb{A}\mathbb{O}_H$ constraint to `cacheAdv`.

To enforce these restrictions, each membrane is created by specifying a tag that represents a control flow behavior, according to Rinar *et al.*'s categories described in Section 4. Thus the cache membrane only accepts aspects with narrowing advice (line 10); and the quota membranes only accept aspects with augmentation advice

⁶The 0rd a constraint is because the keys of a dictionary must be ordered.

(lines 11 and 12). In addition, we use the special `NoAdvice` tag to specific that the browser membrane does not accept aspects (line 9). Concretely, the aspect registration function `register` is overloaded, and the tag is used to choose the appropriate implementation that statically enforces the corresponding control flow interference pattern.

6.2 Data Flow Interference

The components of the system are defined on a constrained but abstract monadic stack. To actually use these components we need to define a concrete monadic stack with the proper requirements. In our example (Figure 4) we use the stack $\mathfrak{M}$ (line 2), which provides effects for aspect weaving with membrane semantics ($\text{AO}_T \text{MB}_T$), write access to a `Log` value ($\mathbb{W}_T \text{Log}$, as required by the quota advice), and read-write access to a `History` and a `Cache` value ($\mathbb{S}_T \text{History}$ and $\mathbb{S}_T \text{Cache}$, as required by the browser and the cache advice).

Interference Problems If all the components share the full monadic stack the cache component cannot work because, due to implicit lifting, the $\mathbb{S}_T \text{Cache}$ layer is hidden. Thus the `get` and `put` operations on the advice will be redirected to the $\mathbb{S}_T \text{History}$ layer, producing a typechecking error. This problem can be solved by using explicit lifting in the cache advice, but this solution breaks the premise of independently developed components, because the cache advice must be aware of other aspects.

A more serious problem is that state layers are accessible to every component. For instance, if the developer adds a $\mathbb{S}_M \text{History}$ $\mathfrak{m}$ constraint to the quota advice, then the history of the browser can be modified by any quota aspect. Eventually any component could write to the log, or even deploy aspects. It could be argued that this is not the case in our example since the constraints on each component are minimal – however the problem is that it is *possible* to use *some* component that interferes with the behavior of the others.

Solution Using Monad Views Our solution is based on the premise that a component must only *see* what it actually requires from the monadic stack – and nothing else. To this end, a membrane receives as a second argument a *monad view* (Section 3.2), which maps between $\mathfrak{M}$ and the virtual stack given to aspects inside the membrane.

A view is constructed using a *structural mask*⁷, where it is specified from left to right whether the respective layer is visible or not. The mask is created by concatenating unary masks for each layer: $\square$ indicates a visible layer, $\blacksquare$ a hidden layer, and $\diamond$ a read-only layer (must originally be a state layer)⁸. For the browser membrane we use the identity view `idv`, which yields $\mathfrak{M}$.

Extending the Pointcut Language Observe that the cache advice is triggered by calls to `browser` detected by a pointcut, and that for type safety reasons the monad stack associated to the pointcut and the advice must be the same [11].

Hence it is not possible to register an aspect that applies the cache advice on calls to the browser using only the predefined `pcCall` pointcut (Section 2), because its type is $(a \rightarrow \mathfrak{m} \ b) \rightarrow \text{PC } \mathfrak{m} \ a \ b$ [11]. That is, the function must see the same monad as the resulting pointcut – but `browser` sees the full stack, while `cacheAdv` sees a restricted monad. Moreover, even if they see the same virtual stack the browser cannot access its required state layer, which is hidden. To solve this issue we provide a new pointcut, `pcCallGlobal`, with type

$$n \bowtie \text{AO}_T \text{MB}_T \ \mathfrak{m} \rightarrow (a \rightarrow \text{AO}_T \text{MB}_T \ \mathfrak{m} \ b) \rightarrow \text{PC } n \ a \ (n \ b)$$

where $\mathfrak{m}$ and n are monads, and $\text{AO}_T \text{MB}_T \ \mathfrak{m}$ is any stack with membrane semantics for weaving. The type $n \bowtie \text{AO}_T \text{MB}_T \ \mathfrak{m}$ denotes any monad view that maps between a full stack $\text{AO}_T \text{MB}_T \ \mathfrak{m}$ and a restricted view n . Therefore with `pcCallGlobal` it is possible to advise calls to functions that see the full monadic stack with advice on a restricted monad.

A `pcCallGlobal` pointcut is non-interfering because it is *pure*, i.e. there are no side-effecting operations on the full stack. This is important because the typing issue could be solved by lifting a `pcCall` pointcut to the proper type using views – but then we could not ensure non-interference in general if the pointcut is composed using pointcut combinators. For instance, one could lift a pointcut that combines `pcCall` with a pointcut that writes to a layer of the full stack that is unavailable in the restricted view, without type errors.

Finally, a more general instance of the same typing issues arises when registering the cache quota aspect, which we solve with the `pcCallView` pointcut, with type

$$1 \bowtie \text{AO}_T \text{MB}_T \ \mathfrak{m} \rightarrow n \bowtie \text{AO}_T \text{MB}_T \ \mathfrak{m} \rightarrow (a \rightarrow 1 \ b) \rightarrow \text{PC } n \ a \ (n \ b)$$

that takes into account a function and two views, 1 and n , which correspond to the view seen by the function, and the view of the resulting pointcut, respectively; and yields a proper pointcut, in the same way as `pcCallGlobal`.

7. Conclusions and Future Work

We have shown a complete example of how to control aspect interference using monadic reasoning about effects, control flow combinators, and membranes. Non-interference is directly enforced by the type system, not by external analysis tools as in other approaches. We believe this system is a straightforward approach to specify and enforce by construction the allowed interactions not only between aspects, but also between aspects and components in a system.

A rigorous formalization of interference scenarios and how they are managed by our system is an immediate line of future work. It would also be interesting to perform some case studies to assess the capabilities of our framework in practice.

Acknowledgements

This work was supported by the INRIA Associated team RAPIDS. We also thank the anonymous reviewers for their valuable recommendations.

References

- [1] *Proceedings of the 9th ACM International Conference on Aspect-Oriented Software Development (AOSD 2010)*, Rennes and Saint Malo, France, Mar. 2010. ACM Press.
- [2] C. Clifton and G. Leavens. Observers and assistants: A proposal for modular aspect-oriented reasoning. In *Proc. of the International Workshop on Foundations of Aspect-Oriented Languages (FOAL 2002)*, 2002.
- [3] C. Disenfeld and S. Katz. A closer look at aspect interference and cooperation. In *Proceedings of the 11th International Conference on Aspect-oriented Software Development (AOSD)*, 2012.
- [4] R. Douence, P. Fradet, and M. Südholt. A framework for the detection and resolution of aspect interactions. In D. Batory, C. Consel, and W. Taha, editors, *Proceedings of the 1st ACM SIGPLAN/SIGSOFT Conference on Generative Programming and Component Engineering (GPCE 2002)*, volume 2487 of *Lecture Notes in Computer Science*, pages 173–188, Pittsburgh, PA, USA, Oct. 2002. Springer-Verlag.
- [5] S. Liang, P. Hudak, and M. Jones. Monad transformers and modular interpreters. In *Proceedings of the 22nd ACM Symposium on Prin-*

⁷There are also *nominal masks* where each layer is labeled with a name. For now we only consider structural masks. For more details see [9].

⁸This graphical representation of views is adapted from the original proposal [9], and is intended to ease the understanding of how masks work. The corresponding code can be found in the implementation.

ciples of Programming Languages (POPL 95), pages 333–343, San Francisco, California, USA, Jan. 1995. ACM Press.

- [6] H. Masuhara, G. Kiczales, and C. Dutchyn. A compilation and optimization model for aspect-oriented programs. In G. Hedin, editor, *Proceedings of Compiler Construction (CC2003)*, volume 2622 of *Lecture Notes in Computer Science*, pages 46–60. Springer-Verlag, 2003.
- [7] B. C. d. S. Oliveira, T. Schrijvers, and W. R. Cook. EffectiveAdvice: disciplined advice with explicit effects. In *AOSD 2010* [1], pages 109–120.
- [8] M. Rinard, A. Salcianu, and S. Bugrara. A classification system and analysis for aspect-oriented programs. In *Proceedings of the 12th ACM Symposium on Foundations of Software Engineering (FSE 12)*, pages 147–158. ACM Press, 2004.
- [9] T. Schrijvers and B. C. Oliveira. Monads, zippers and views: virtualizing the monad stack. In *Proceedings of the 16th ACM SIGPLAN Conference on Functional Programming (ICFP 2011)*, pages 32–44, Tokyo, Japan, Sept. 2011. ACM Press.
- [10] M. Störzer, R. Sterr, and F. Forster. Detecting precedence-related advice interference. In *Proc. of the 21st International Conference on Automated Software Engineering (ASE 2006)*, 2006.
- [11] N. Tabareau, I. Figueroa, and É. Tanter. A typed monadic embedding of aspects. In J. Kinzle, editor, *Proceedings of the 12th International Conference on Aspect-Oriented Software Development (AOSD 2013)*, Fukuoka, Japan, Mar. 2013. ACM Press.
- [12] É. Tanter. Execution levels for aspect-oriented programming. In *AOSD 2010* [1], pages 37–48.
- [13] É. Tanter, P. Moret, W. Binder, and D. Ansaloni. Composition of dynamic analysis aspects. In *Proceedings of the 9th ACM SIGPLAN International Conference on Generative Programming and Component Engineering (GPCE 2010)*, pages 113–122, Eindhoven, The Netherlands, Oct. 2010. ACM Press.
- [14] É. Tanter, N. Tabareau, and R. Douence. Taming aspects with membranes. In *Proceedings of the 11th Workshop on Foundations of Aspect-Oriented Languages (FOAL 2012)*, pages 3–8, Potsdam, Germany, Mar. 2012. ACM Press.
- [15] K. Tian, K. Cooper, K. Zhang, and S. Liu. Towards a new understanding of advice interference. In *Proceedings of the 4th International Conference on Secure Software Integration and Reliability Improvement (SSIRI)*, 2010.
- [16] K. Tian, K. Cooper, K. Zhang, and H. Yu. A classification of aspect composition problems. In *Proceedings of the 3th International Conference on Secure Software Integration and Reliability Improvement (SSIRI)*, 2009.