Towards an Aspect-oriented Language Module: Aspects for Petri Nets

Tim Molderez * Bart Meyers Dirk Janssens Hans Vangheluwe

Dept. of Mathematics and Computer Science
University of Antwerp, Belgium
{tim.molderez,bart.meyers,dirk.janssens,hans.vangheluwe}@ua.ac.be

Abstract

The concept of composing a (domain-specific) language from different reusable modules has gained much interest over the years. The addition of aspect-oriented features to a language is a suitable candidate of such a module. However, rather than directly attempting to design an aspect-oriented language module that is applicable to any base language, this paper focuses on adding aspect-oriented features to a language that is quite different from prevalent base languages (e.g. Java): Petri nets. A running example demonstrates the use of aspects to enforce an invariant on a base Petri net.

Categories and Subject Descriptors D.3.3 [*Programming Languages*]: Language Constructs and Features; I.6.5 [*Simulation and Modelling*]: Model Development

General Terms Languages, Design

Keywords Petri nets, aspect-oriented modelling, language engineering

1. Introduction

Creating a new (domain-specific) language, including the tools, documentation and community that accompany a language, is an undeniably large endeavour. To reduce this effort, there is a growing interest in constructing new languages by extending existing languages, or by composing different languages in a modular manner [2, 7, 8]. Next to this research area, there also is a large body of work that present aspect-oriented versions of various different languages. The combination of these two research areas leads to the notion of an "aspect-oriented language module". In other words, a module that implements (pointcut-and-advice flavoured) aspect-oriented features, and is applicable to a wide range of base languages. However, rather than attempting to cover all base languages at once, this paper focuses on adding aspect-oriented features to just one language: Petri nets. We chose the Petri nets modelling language as our base language, simply because it is so different from commonly used base languages (e.g. Java): Petri nets are non-deterministic, their state is implicit and the language is very minimal. In this paper, we are mainly interested in how this choice of base language affects the design of aspect-oriented features. Consequently, we present an aspect-oriented extension to Petri nets and demonstrate its use with an invariant enforcement example. While designing this aspect-oriented Petri nets extension, we keep the general idea of an aspect-oriented language module in mind. This is done by dividing the extension itself into a number of components, such that the extension is no longer specific to Petri nets at this level of abstraction.

The remainder of this paper is structured as follows: Sec. 2 introduces the paper's running example. Sec. 3 then elaborates on this example while the various components of the aspect-oriented extension to Petri nets are discussed. Sec. 4 presents related work; Sec. 5 concludes the paper and discusses future work.

2. Invariant enforcement example

We will use a running example to demonstrate the use of our aspect-oriented Petri net extension. The "base Petri net", i.e. the Petri net without any aspects, for this example is shown in Fig. 1. The places within the dotted rectangle (`P5-P8`) represent a number of rooms within a building. All places outside this rectangle (`P1-P4`) represent the outside world. The transitions going in and out of the rectangle represent the building's various entrances (`T1,T3`) and exits (`T2,T4`). The goal of this example is to enforce an invariant: Only a certain number of people (tokens) may be inside the building at the same time. If more people were to enter the building, this could, for example, violate the building's fire safety

* Funded by a doctoral scholarship of the Research Foundation - Flanders (FWO)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

DSAL '12, March 27, 2012, Potsdam, Germany.

Copyright © 2012 ACM 978-1-4503-1128-1/12/03. . . $10.00

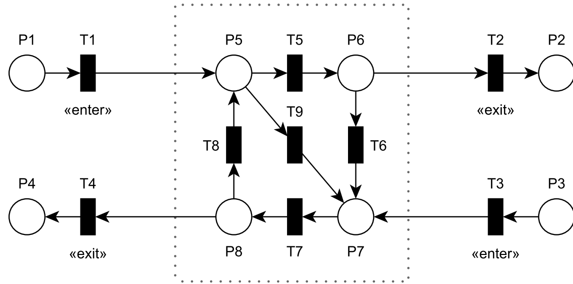

Figure 1: Base Petri net

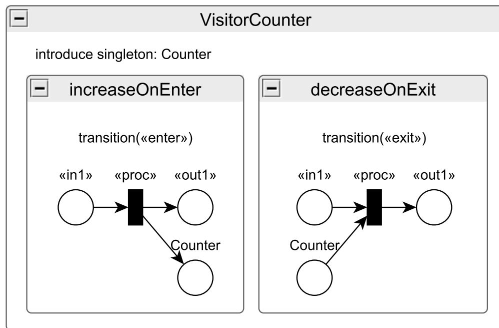

Figure 2: `VisitorCounter` aspect

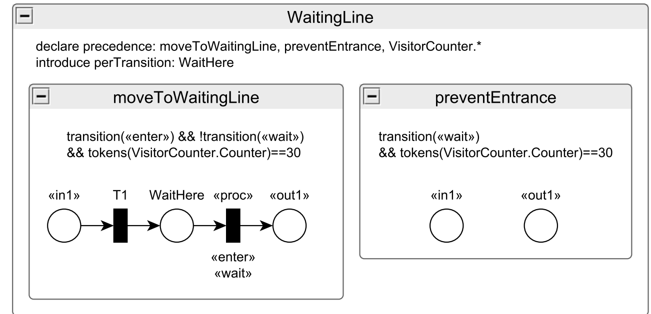

Figure 3: `WaitingLine` aspect

regulations. To prevent this from happening, the following sections will use aspects to monitor the building's entrances and exits.

The aspects enforcing the invariant are shown in Fig. 2 and Fig. 3. The first aspect, called `VisitorCounter`, keeps track of how many people are currently inside the building. It contains two pointcut-advice pairs: The pair called `increaseOnEnter` increases the `Counter` place whenever someone passes through an entrance. That is, its pointcut captures all occurrences of transitions annotated with an «`enter`» stereotype. While our pointcuts make use of these stereotypes, note that stereotypes are not strictly necessary. We could have also explicitly listed the names of the transitions that should be matched by the pointcut. However, the use of stereotypes makes the pointcuts easier to understand and less dependent on concrete transition names, therefore reducing the impact of the fragile pointcut problem. This is based on Noguera et al. [6], which uses Java annotations within AspectJ pointcuts.

As expected, the `decreaseOnExit` pointcut-advice pair decreases the `Counter` place whenever someone leaves the building, i.e. an «`exit`» transition is fired.

The second aspect, `WaitingLine`, enforces the invariant that only a certain number of people, say 30, may be in the building at the same time. This aspect is dependent on the `VisitorCounter` aspect, as it needs access to its `Counter` place. The `moveToWaitingLine` pointcut-advice pair acts whenever the building is full and someone tries to enter. That person will then be redirected to a waiting line, represented by the `WaitHere` place. There is such a waiting line for each entrance. If someone in the waiting line tries to enter the building while it still is full, the `preventEntrance` pair will prevent that person from entering. To make sure that the `Counter` place is only increased if the building is not full, the `declare precedence` statement expresses that the pointcuts in the `WaitingLine` aspect have precedence over those in `VisitorCounter`. We will elaborate on the semantics of these aspects over the course of the following sections.

3. Overview of the aspect-oriented extension

To establish aspect-oriented Petri nets as an extension to Petri nets, there are a number of different components to be determined: the base language (Petri nets), a weaver, join point model, pointcut language, advice language and a composition mechanism. Note that, at this level, none of these components are specific to Petri nets; they may as well be applied to another base language.

3.1 Weaver

The core component of the extension is the weaver, which manages most communication between the other components. It roughly performs the following steps:

1. During execution, whenever a join point (as determined by the join point model of Sec. 3.2) is reached, normal Petri net execution is paused and control is passed to the weaver.
2. Given the current state and the current join point, the weaver tests all pointcuts to find the ones that match.
3. Given the pointcuts that match, the composition mechanism determines the right order in which the corresponding advice should be inserted.
4. Each advice now is inserted into the base Petri net to create an augmented Petri net. This is done in the order determined by the composition mechanism.
5. Petri net execution resumes.

3.2 Join point model

In general, any point in time during a model/program execution could be considered an element of a language's join point model. In case of Petri nets, we have chosen our join point model to consist of all run-time occurrences of transitions. That is, advice can be executed whenever a transition is *about to* fire. We have chosen this particular join point model, because it is not affected by the non-deterministic nature of Petri nets: Whenever a Petri net execution arrives at a particular join point, it has not just determined which transitions *can* be fired (i.e. which transitions are enabled); it has also chosen which transition(s) *will* be fired. Therefore, it is clear which behaviour will be extended or replaced whenever an advice is applied.

3.3 Pointcut language

After choosing a join point model, we can create the language used to describe pointcuts. The language that we have chosen for aspect-oriented Petri nets is minimal: Our pointcuts are boolean expressions that can make use of two constructs: `transition` and `tokens`. An example of `transition` is shown in Fig. 2: `transition(«enter»)`. This construct is true if the current join point, being an occurrence of a transition, corresponds to one of the parameters in the `transition` construct. In this example, the parameter indicates all transitions with the `«enter»` stereotype. It is of course also possible to pass in concrete transition names as parameters, or to make use of wildcards. Note that the `transition` construct does not make use of the join point model's dynamic nature, as the mapping to join point shadows, i.e. transitions (not *occurrences of* transitions), is trivial. Pointcuts get more interesting when combined with the `tokens` construct. It is used to reason about the amount of tokens in a place. We have used it in the pointcuts of the `WaitingLine` aspect in Fig. 3. For example, `tokens(VisitorCounter.Counter)==30` is true if there are 30 tokens in the `Counter` place of the `VisitorCounter` aspect.

3.4 Advice language

If a pointcut matches in a pointcut-advice pair, the corresponding advice should be inserted into the base Petri net, resulting in an augmented Petri net. Note that we do not make a distinction between before, after and around advice. As this distinction would require additional syntax, and because before and after advice are only specialized cases of around advice, we simply consider all advice to be around advice. In our extension, an advice looks like a regular Petri net, but is extended with a few different constructs. First, there are special stereotypes for places that are used to specify how the input and output places of a join point should be bound to an advice. This is similar to the notion of ports in hierarchical Petri nets [4]. These bindings are discussed in Sec. 3.4.1. Second, transitions can also carry a `«proc»` stereotype, indicating a proceed transition, similar in concept to proceed calls. Proceed transitions are covered in Sec. 3.4.2. Finally, places can be shared among advice by means of introductions (also known as inter-type declarations in AspectJ), introduced in Sec. 3.4.3.

3.4.1 Input and output place binding

Looking at the advice of `increaseOnEnter` in Fig. 2, the `«in1»` and `«out1»` stereotypes instruct how this advice must be bound to the input and output places of a transition. In this case, the `transition(«enter»)` pointcut always matches whenever a transition with an stereotype `«enter»` is about to be fired. Note that these `«enter»` transitions always have one input and one output place. When applying our advice, what we want to achieve is to intercept the firing of an `«enter»` transition and to replace it with an advice execution. To be able to do this, we need to bind the intercepted `«enter»` transition's input place to the advice's `«in1»` place. Similarly, the output place should be bound to the `«out1»` place. Note that the number of input and output places of the intercepted transition must match with the `«inN»` and `«outN»` stereotypes in the advice. For example, if the `«enter»` transition in the base net would have two input places, our advice must also have an `«in2»` place. Otherwise, the advice cannot be applied, which implies that the pointcut does not match. In this sense, the `«inN»` and `«outN»` stereotypes also form a part of the pointcut language. Additionally, if an advice can be applied and there are multiple input or output places to be bound, the choice of which input/output place must bind to which `«inN» / «outN»` place is currently left non-deterministic.

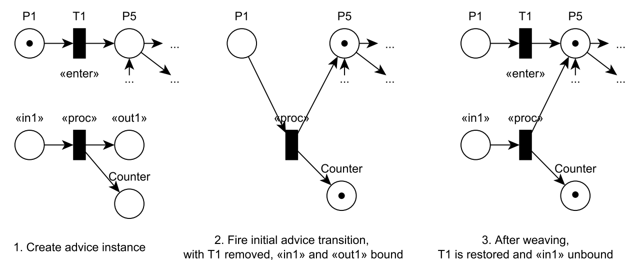

Figure 4: Weaving an advice into the base Petri net

After introducing the `«inN»` and `«outN»` stereotypes, inserting an advice into a base Petri net is performed as follows: Suppose that `increaseOnEnter`'s pointcut just matched on transition `T1` in the base Petri net and we wish to insert the corresponding advice. We will also assume that only one pointcut matched on this join point, in which case the proceed transition (with a `«proc»` stereotype) acts as a normal transition. (Shared join points are discussed in Sec. 3.4.2.)

To insert `increaseOnEnter`'s advice, a new instance of the advice is created and added to the base Petri net. This is shown in step 1 of Fig. 4. (Only the relevant part of the base Petri net is shown.) This advice instance is created such that the `«out1»` place is bound to the join point's output

place, P5. Note that this binding is permanent. The input place's binding only holds for an instant however: The effect that we want to achieve is, instead of firing T1, we want to fire the transition attached to the advice's «in1» place. To accomplish this effect, we temporarily remove T1 and bind the input place, i.e. «in1» is bound to P1. This only lasts until we have fired. This situation is shown as step 2 in Fig. 4. Once this is done, T1 is added to the net again, «in1» is no longer bound and, as of now, is nothing but a regular (empty) place. At this point, we have a normal Petri net once more, meaning that advice insertion has completed and regular Petri net execution can continue. This resulting Petri net is shown as step 3 in Fig. 4.

The reason that the input place's binding is removed as soon as the advice initially fired, is because our advice insertion should only intercept one single join point, i.e. one occurrence of a transition. Otherwise the advice would interfere with subsequent occurrences of that transition. On the other hand, the output place binding is permanent because we wish to preserve non-determinism: We do not enforce the advice to finish its execution (i.e. to continue firing transitions in the advice as long as possible) before execution of the base Petri net can resume. As the resulting Petri net is a normal Petri net, it is allowed that transitions from the base net are fired even though the advice has not finished its execution yet. To allow the result of the advice execution to be passed to the base net, binding of output places is permanent.

Another design choice worth discussing is why we create a new advice instance for *every* matching join point. If one aspect applies to multiple join points, its instance may not be reusable because the bindings of its proceed transition(s), which will be discussed in Sec. 3.4.2, could be different per join point. The main disadvantage of creating a new instance for each matching join point is of course that the augmented Petri net easily grows in size. However, this can be heavily optimized: For example, if «outN» places only have incoming edges, then an advice instance can be safely removed as soon as no more enabled transitions are found within the instance. Such an instance can only produce output to the base net, and if no transitions are enabled, nothing can be produced and therefore the instance can be removed without altering the augmented Petri net's behaviour.

3.4.2 Proceed transitions

As mentioned earlier, a transition with a «proc» stereotype indicates a proceed transition. Similar to AspectJ, an advice usually contains one proceed transition, but it may contain none or multiple of them as well. If there are multiple advice at the same join point, the proceed transitions indicate where the next advice in the aspect composition must be woven. In order to make this possible a proceed transition must also be compatible with the number of input and output places of the join point that was intercepted. For example, if a pointcut captures transitions with one input and one output place, a proceed transition must (at least) have one input and one output place as well. (In case ambiguities arise, this can be resolved by adding stereotypes to map the input and output places of the proceed transition.) Our invariant enforcement example also includes an instance where join points are shared by multiple pointcuts: Once the building is at full capacity, i.e. Counter contains 30 tokens, then the pointcuts of moveToWaitingLine, preventEntrance and increaseOnEnter will match. As determined by the composition mechanism, we first create an instance of moveToWaitingLine. Immediately after creating this instance, the preventEntrance advice is inserted into the proceed transition of moveToWaitingLine. Note that «inN» and «outN» bindings are permanent when inserting into a proceed transition. If the preventEntrance advice would contain a proceed transition, increaseOnEnter would now be inserted there. The advice preventEntrance however consciously does not contain a proceed transition, as its purpose is to prevent entrance to the building. Once all proceed transitions are resolved, we can continue the insertion process of moveToWaitingLine; i.e. bind its inputs and outputs, initiate advice execution and remove the input binding.

3.4.3 Introductions

To be able to share information between different advice instances, introductions can be used. This is used in the VisitorCounter aspect in Fig. 2 to make the Counter place global. The introduce singleton: Counter statement in the aspect declares that there only is one instance of Counter, which is shared among all advice instances in VisitorCounter. That is, if there is a place named Counter in the advice, that place will be bound to the actual Counter instance whenever an advice instance needs to be inserted. If there are multiple instances of increaseOnEnter or decreaseOnExit, the Counter place will also get several incoming and outgoing edges.

The WaitingLine aspect in Fig. 3 makes use of introductions as well. Rather than introducing a global place, it introduces WaitHere places with the introduce perTransition statement. This statement declares that a WaitHere place will be created for each transition corresponding to the join points that matched. In other words, all advice that matched on the same join point shadow will share the same WaitHere place. In the context of our example, a WaitHere place is thus created for all entrances to the building, as moveToWaitingLine's pointcut looks for «enter» transitions.

3.5 Composition mechanism

The final component of the aspect-oriented Petri net extension is the composition mechanism. Similar to AspectJ, we use a declare precedence statement for this mechanism. An example is shown in the WaitingLine aspect of Fig. 3. Our mechanism is slightly more fine-grained, as the composition order is defined at the level of pointcuts instead of aspects. This level of granularity is needed for our example,

as people should be redirected to a `WaitHere` place before preventing access to the building, in case it is at full capacity. In other words, `moveToWaitingLine` should always be executed before `preventEntrance`. Additionally, these two must be executed before `increaseOnEnter`; otherwise the `Counter` place could increase even if this is not allowed.

4. Related work

In terms of related work, the closest to our aspect-oriented Petri nets extension is the work presented in Xu et al. [9], where aspects are used in Petri nets to implement threat mitigations in security design. Their join point model however only focuses on structural join points, consisting of places and transitions. As a result, the weaving process is implemented as a preprocessing step without any residual logic. While easier to understand, our interests lie more in exploring how characteristic aspect-oriented constructs carry over to different base languages. The feature-oriented Petri nets extension in Muschevechi et al. [5] is used to model software product lines. Its extension to Petri nets adds transitions that are guarded (also called application conditions) by which features are selected in a software product line. The approach is only tested for small examples, so it is currently unclear how the approach scales to larger systems, especially when crosscutting features are involved. It may be worthwhile to use aspect-oriented techniques to support featured-oriented concepts, as this has been done before at the level of programming languages in Apel et al. [1]. The work of Hamadi et al. [3] presents a domain-specific Petri nets extension to specify exceptional behaviour in workflow systems. A net is composed of a number of regions that can be dynamically restructured based on execution history. Aspect-oriented Petri nets may as well be suited for this purpose, as pointcuts and advice also allow for dynamic Petri net modifications. To do this in terms of execution history, support for tracematches should be added. Finally, there also is a close connection with hierarchical Petri nets [4], as the use of aspects automatically introduces a notion of modularity and reuse. The main difference is that aspects implement inversion of control and are more implicit, whereas hierarchical Petri nets are explicit. In combination with pointcuts, this makes aspects better suited for modularising crosscutting concerns. Additionally, introductions make it possible to share information between aspects and the base net.

5. Conclusion and future work

This paper has presented an aspect-oriented extension to Petri nets, as an initial step to an aspect-oriented language module. The Petri nets language was chosen because it distinguished itself from common base languages through its non-determinism, implicit state and its simplicity. The running example has shown an interesting use for aspect-oriented Petri nets: Aspects can be used to enforce invariants in a compact and modular manner. The resulting Petri net may then be used for further analysis. One interesting direction of future work however is to study how our addition of aspects affects its use for analysis. For example, if a base Petri net is shown to be free of deadlocks, what happens if an aspect is added? Which characteristics should this aspect have in order to preserve this desired property? Additionally, because Petri nets are such a small language compared to most programming languages, aspect-oriented Petri nets may also be an interesting use case to study the interactions among aspects, and between aspects and the base system. An implementation of the language using graph transformation rules should be fairly straightforward, and can be suited to explore these interactions. Another direction of future work is more towards the general-purpose aspect-oriented language module: A variety of other base languages can be extended with aspect-oriented features to form a more precise idea of what it means for a language to be aspect-oriented, and what this language module may look like.

References

[1] Sven Apel, Thomas Leich, Marko Rosenmüller, and Gunter Saake. FeatureC++: on the symbiosis of Feature-Oriented and Aspect-Oriented programming. In *Generative Programming and Component Engineering*, volume 3676, pages 125–140. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005.

[2] S. Dmitriev. Language oriented programming: The next programming paradigm. *JetBrains onBoard*, 1(2), 2004.

[3] Rachid Hamadi and Boualem Benatallah. Dynamic restructuring of recovery nets. In *Proceedings of the 16th Australasian database conference - Volume 39*, ADC '05, page 37–46, Darlinghurst, Australia, 2005.

[4] P. Huber, K. Jensen, and R. Shapiro. Hierarchies in coloured petri nets. *Advances in Petri Nets 1990*, page 313–341, 1991.

[5] R. Muschevici, D. Clarke, and J. Proenca. Feature petri nets. In *Proceedings of the 14th International Software Product Line Conference (SPLC 2010)*, volume 2, 2010.

[6] Carlos Noguera, Andy Kellens, Dirk Deridder, and Theo D'Hondt. Tackling pointcut fragility with dynamic annotations. In *Proceedings of the 7th Workshop on Reflection, AOP and Meta-Data for Software Evolution*, RAM-SE '10, page 1:1–1:6, New York, NY, USA, 2010. ACM.

[7] Charles Simonyi, Magnus Christerson, and Shane Clifford. Intentional software. In *Proceedings of the 21st annual ACM SIGPLAN conference on Object-oriented programming systems, languages, and applications*, OOPSLA '06, page 451–464, New York, NY, USA, 2006. ACM.

[8] Markus Völter and Eelco Visser. Language extension and composition with language workbenches. In *Proceedings of the ACM international conference companion on Object oriented programming systems languages and applications companion*, SPLASH '10, New York, NY, USA, 2010. ACM.

[9] D. Xu and K. E Nygard. Threat-driven modeling and verification of secure software using aspect-oriented petri nets. *IEEE Transactions on Software Engineering*, 32(4), 2006.