
Model-driven and pattern-based integration of process-driven SOA models

Uwe Zdun* and Schahram Dustdar

Distributed Systems Group,
Information Systems Institute,
Argentinierstrasse 8/184-1,
Wien A-1040, Austria
E-mail: zdun@infosys.tuwien.ac.at
E-mail: dustdar@infosys.tuwien.ac.at
*Corresponding author

Abstract: Service-Oriented Architectures (SOA) are increasingly used in the context of business processes. However, the modelling approaches for process-driven SOAs do not yet sufficiently integrate the various kinds of models relevant for a process-driven SOA – ranging from process models to software architectural models to software design models. We propose to integrate process-driven SOA models via a model-driven software development approach that is based on proven practices documented as software patterns. We introduce pattern primitives as an intermediate abstraction to precisely model the participants in the solutions that patterns convey. To enable model-driven development, we develop domain-specific modelling languages for each kind of process-driven SOA model – based on meta-models that are extended with the pattern primitives. The various process-driven SOA models are integrated in a model-driven tool chain via the meta-models. Our tool chain validates the process-driven SOA models with regard to the constraints given by the meta-models and primitives.

Keywords: SOA; process-driven SOA; software patterns; services modelling.

Reference to this paper should be made as follows: Zdun, U. and Dustdar, S. (2007) 'Model-driven and pattern-based integration of process-driven SOA models', *Int. J. Business Process Integration and Management*, Vol. 2, No. 2, pp.109–119.

Biographical notes: Uwe Zdun is an Assistant Professor at the Vienna University of Technology. His research interests include software patterns, software architecture, SOA, distributed systems, language engineering and object orientation. He received his Doctoral Degree in Computer Science from the University of Essen in 2002 and his Habilitation degree (*venia docendi*) from Vienna University of Economics and BA in 2006. He is the coauthor of the books *Remoting Patterns* (John Wiley & Sons, 2004) and *Software-Architektur* (Elsevier/Spektrum, 2005).

Schahram Dustdar is a Full Professor for the Internet Technologies at the Distributed Systems Group, Information Systems Institute, Vienna University of Technology (TU Wien) where he is the Director of the Vita Lab and an Honorary Professor of Information Systems at the Department of Computing Science at the University of Groningen (RuG), the Netherlands. He received his MSc (1990) and PhD (1992) in Business Informatics (Wirtschaftsinformatik) from the University of Linz, Austria. He received his Habilitation degree (*venia docendi*) in April 2003.

1 Introduction

Many Service-Oriented Architectures (SOA) provide a service composition layer that introduces a process engine as the top-level layer (cf. Zdun et al., 2006). Services realise individual activities in the process. This kind of architecture is called *process-driven SOA*. The main goal of process-driven SOAs is to increase the productivity, efficiency and flexibility of an organisation. This is achieved by aligning the high-level business processes with the technical IT services. That is, the business goals get closely integrated with the IT architecture.

One of the most important characteristics of SOAs suggests heterogeneity of technologies and integration across

vendor-specific technologies (cf. Vinoski, 2003). This, however, yields an important challenge for modelling process-driven SOAs: many modelling domains need to be considered and the different kinds of models need to be integrated. Among many other modelling domains, we need to consider component architectures, message flows, transactions, security, workflows/business processes, programming language (snippets), business object designs and organisational models. Application domains additionally introduce domain models, such as banking or insurance domain models. Also, implicit or explicit models for integrating existing legacy systems are needed, sometimes using different concepts than the rest of the SOA.

In other words, a central challenge for modelling process-driven SOAs is that we generally need to integrate different kinds of models and abstractions. This problem is challenging because so far there is no modelling approach for integrating all these kinds of models.

In this paper, we propose a concept for a model-driven tool chain that addresses these challenges through Model-Driven Software Development (MDS) (cf. Greenfield and Short, 2004; Stahl and Voelter, 2006). Our concept is based on the precise specification of the models in Domain-Specific Languages (DSL) which are themselves precisely specified by meta-models and constraints on them. The code in the executable languages is generated from the models expressed in the DSLs. That is, the integration issues raised above are solved at the meta-model level.

Our tool chain and MDS concepts break the integration issues down to the problem of finding adequate meta-models for representing all concerns to be modelled in the various modelling languages used in the modelling domains. In this paper, we propose to develop the meta-models according to proven practices that can be found in existing process-driven SOAs. Our assumption is that using proven practices as a foundation for meta-modelling leads to a close match between the modelling abstractions and the real-world requirements.

In our approach, software patterns are used to describe the proven practices. Software patterns capture reusable design knowledge and expertise that provides proven solutions to recurring software design problems that arise in particular contexts and domains (cf. Schmidt and Buschmann, 2003). A software pattern, however, is described in an informal form and cannot easily be described formally, for example, by using a parameterisable, template-style description. Hence, as such, patterns are not usable as elements of meta-models. We remedy this problem by introducing an intermediate abstraction called pattern primitives. A pattern primitive is a fundamental, precisely specified modelling element in representing a pattern.

Our general approach to apply pattern primitives for process-driven SOAs is to define meta-models for all kinds of models that are needed and specify the pattern primitives as extensions of these meta-models. The connection between the various kinds of models and the validation of the models – with regard to model integrability and consistency in and across the modelling domains – is the main task of our model-driven tool chain concepts. To demonstrate our approach, we use a precisely specified subset of UML2 and OCL to depict the various refinements of processes in process-driven SOA models. We use the precisely specified subset of UML2 and OCL only for demonstration purposes. Any other precisely specified meta-model can be used as well. ‘Precisely specified’ means in this context that no informal elements are part of the meta-models or constraints (like the informal constraints used in some parts of the UML standard) and the semantics of the meta-model or constraint elements are specified in the generator – to define precisely how models need to be validated and how valid code can be generated.

In this paper, we first provide the background on MDS and patterns/pattern primitives in Section 2. Next, in Section 3, we explain our model-driven tool chain to give an overview of our approach. In Section 4, we explain

how we use meta-models to integrate process-driven SOA models. In Section 5, we explain the pattern primitives approach for process-driven integration of services using flow abstractions as the primary modelling domain. In Section 6, we demonstrate how architectural abstractions – as one example of another modelling domain – can be integrated with the flow models. We explain all primitive models with running examples from a pattern language for process-driven integration of services, which we have implemented in our MDS tool chain to validate our approach.

2 Background on MDS

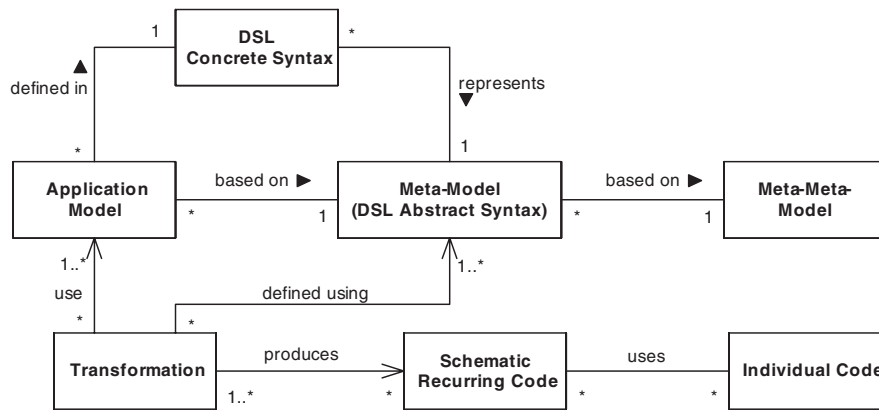
2.1 Model-driven software development

Our approach to MDS (cf. Greenfield and Short, 2004; Stahl and Voelter, 2006)¹ for process-driven SOAs is based on the notion of Domain-Specific Languages (DSL) for modelling the various types of models. DSL are ‘small’ languages that are tailored to be particularly expressive in a certain problem domain. The DSL describes knowledge via a graphical or textual syntax (the DSL’s concrete syntax in the terminology of Greenfield and Short (2004)), which is tied to domain-specific modelling elements through a precisely specified language meta-model (the DSL’s abstract syntax in the terminology of Greenfield and Short (2004)). The meta-model can be instantiated in concrete application models. The semantics of the DSL are defined via unambiguous specifications of model-to-model or model-to-code transformations. There are different ways to specify transformations, such as transformation rules, imperative transformations or template-based transformations.

The meta-models presented in this paper are based on the UML2 meta-model (and extensions of it): for example we use UML2 activity diagrams to model flow abstractions and UML2 class/component diagrams to model object-oriented design and architecture models. To make the UML2 meta-models usable in our approach, we first specify a precise subset of the UML meta-model (using model elements, transformations and constraints). We use the UML only for demonstration purposes (and because it was required in some of our projects); any other meta-models can be used equally. For instance, if a project goal is to generate BPEL code using the full BPEL specification, it is advisable to use a DSL that is closer to BPELs constructs to minimise potential transformation problems.

Meta-models are defined in terms of a meta-meta-model. In UML (and the OMG’s MDA proposal), for instance, this is MOF. Most MDS tools support their own meta-meta-model, which basically represents a mapping from meta-model definitions to the implementation of the MDS tool chain. Below, in the examples from our prototype, we use a simple meta-meta-model to define both the UML2 meta-model and pattern primitives extensions. It is not particularly important for our approach, which meta-meta-model is used, there must be some way to just specify the relationships and transformations between a meta-model and the target models/code of the transformations. The meta-meta-model is not visible to

Figure 1 Relations of artifacts in MDS



developers who build application models, but only to those who build meta-models.

The ultimate goal of the transformations in MDS tools is to generate code in executable languages, such as programming languages or process execution languages. The MDS tools are used to generate all those parts of the executable code which are schematic and recurring and hence can be automated. Of course, some code must be hand-written either because it is individual code for a system or the semantics of the code are not fully covered by the DSLs (yet). The individual code and the generated code use each other and interact through well defined interfaces.

Figure 1 summarises the relations of artifacts in MDS.

2.2 Software patterns and pattern primitives

Software patterns and pattern languages have gained wide acceptance in the field of software development, because they provide systematic reuse strategies for design knowledge (cf. Schmidt and Buschmann, 2003). A pattern encodes proven practice in the form of a reusable design solution to a recurring design problem. A pattern language is a collection of patterns that solves the prevalent problems in a particular domain and context. Patterns informally describe many possible variants of one software solution that a human developer or designer can recognise as one and the same solution.

Even though these properties of the pattern approach are highly valuable in the software design process, they also make pattern instances hard to trace in the models and

implementations. To overcome this problem, we introduced an approach to document precisely specified primitive abstractions that can be found in the patterns (cf. Zdun and Avgeriou, 2005). Documenting pattern primitives means to find precisely describable modelling elements that are primitive in the sense that they represent basic units of abstraction in the domain of the pattern. Our original pattern primitives concept presented in Zdun and Avgeriou (2005) is only targeted at modelling architectural patterns. In this realm, basic architectural abstractions like components, connectors, ports and interfaces are used. An interesting challenge in describing the pattern primitives for the patterns of process-driven SOA is that this area requires various design and architecture concepts, as well as various design and implementation languages.

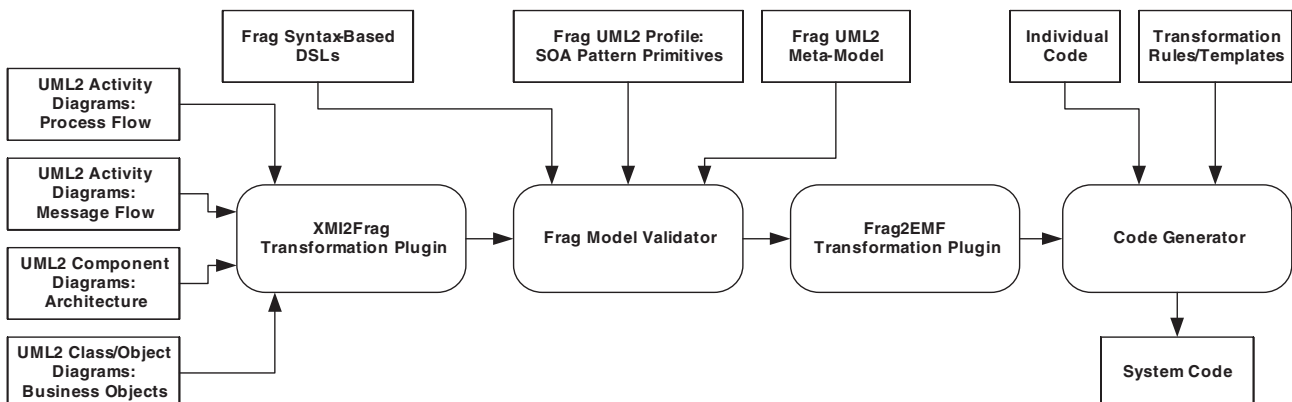
We specify an extension of a meta-class for each elicited primitive, using the UML’s extension mechanisms: stereotypes, tag definitions and constraints. We use the Object Constraint Language (OCL) to precisely specify the constraints of the primitives.

3 MDS for process-driven SOA

3.1 Model-driven tool chain

Our approach requires a model-driven tool chain that loosely follows an architecture similar to our tool chain depicted in Figure 2. We mainly use UML2 models that are extended with UML2 profiles for modelling the pattern primitives as inputs. These UML2 models can either be developed with

Figure 2 Tool chain overview



UML tools (with XMI export) or directly in the textual DSL syntax. If a UML tool is used, the XMI export is transformed into the textual DSL syntax.

We use Frag (cf. Zdun, 2005, 2006) as the syntactic foundation of the textual DSLs and for defining the meta-models of the DSLs. Frag's main goal is to provide a tailorable language. Among other things, Frag supports the tailoring of its object system and the extension with new language elements. Hence, Frag provides a good basis for defining a UML2-based textual DSL because it is easy to define a meta-meta model on top of which we can define the UML meta-classes. Frag automatically provides us with a syntax for defining application models using the UML2 meta-classes. In addition, we have defined a constraint language which follows the OCL's constructs.

The model validator gets all input models and validates the conformance of the application models to the meta-models. It also checks all OCL constraints, in particular, the constraints given by the pattern primitive definitions.

After the model is validated it is transformed into an EMF model, which is understood by the code generator. We then generate code in executable languages using transformation specifications provided to the code generator.

3.2 Model integration concepts: meta-meta-model based integration

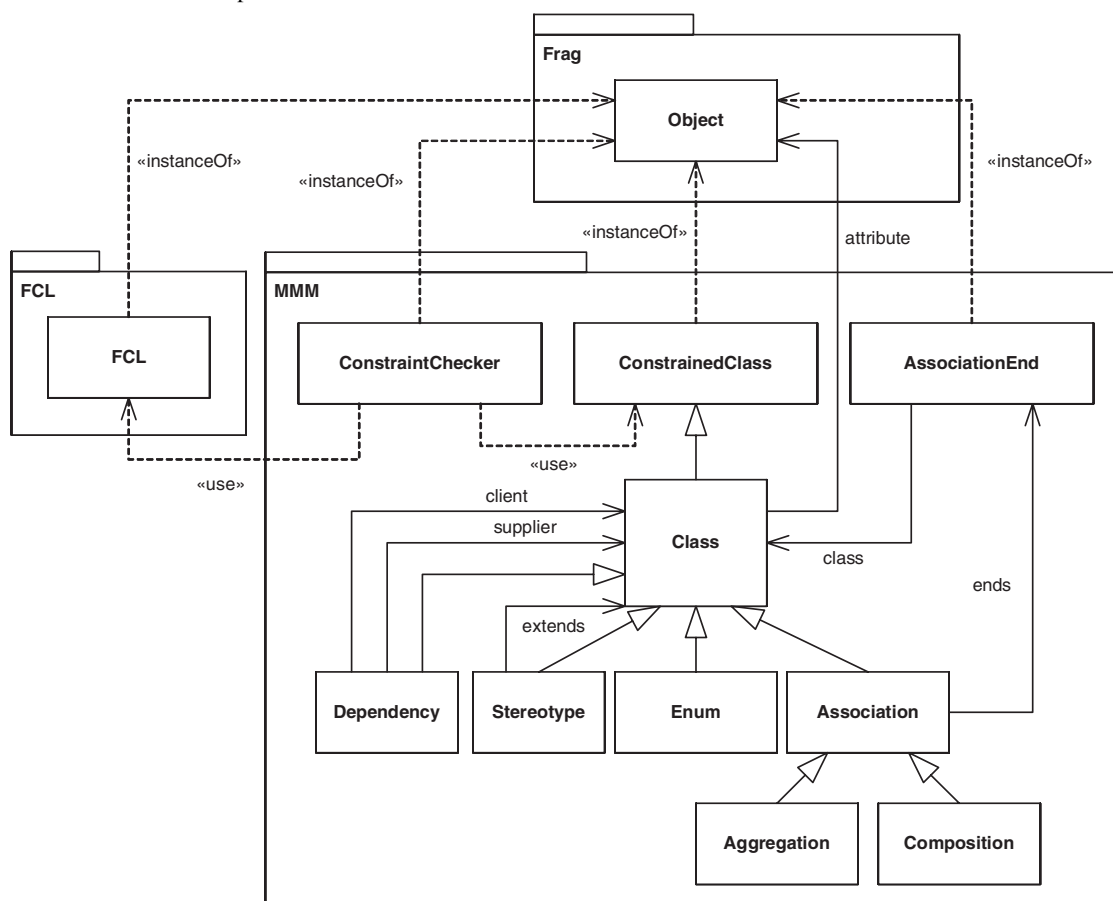
The model-driven concepts described in the previous section only concentrate on the individual modelling domains. For integration of the models, we propose further integration

concepts that extend the general model-driven approach. Because they are independent of external tools, languages or models, in our concept, the central point of integration are the meta-models that we need to define for the DSLs. Also, they are located at the central place of the model-driven architecture: at the point in the tool chain where all different models are assembled.

We propose to define the meta-models on top of one common meta-meta-model. The meta-meta-model can be very simple or more elaborate like MOF. The most important criterion for the meta-meta-model is that the elements of the meta-meta-model allow the model validator to check models against the meta-models. In addition, it should be possible to define a constraint language using the meta-meta-model, with which models can be constrained at the meta-level and hence validated at the model level.

As an example, Figure 3 shows the relevant excerpt of the meta-meta-model that we use in Frag to define precisely specified subsets of the UML2 meta-model. This meta-meta-model is very simple and reuses Frag's language features wherever possible. It is derived from the most general class in the Frag object system: `Object`. The meta-meta-model classes are subclasses of `ConstrainedClass` which allows us to add OCL-style constraints to classes. The convenience class `ConstraintChecker` looks up to all `ConstrainedClass` instances via reflection and checks the constraints. Constraints are specified in a language similar to OCL (defined using the class `FCL`). The meta-models are defined using `Class`. We introduce also a number of relationships between classes: dependencies,

Figure 3 Meta-meta-model excerpt



associations, compositions, and aggregations. In addition, typed attributes can be specified. Please note that we do not define the generalisation relationship, because multiple inheritance is suitably predefined by Frag and we can reuse this implementation. The Stereotype class defines the UML2 extends-relationship, which is used to extend meta-classes. Enum is a convenience class to define Enumeration types.

3.3 Model integration concepts: proven practices-based integration

Besides the common meta-meta-model concept, we use proven practices descriptions as the second central model integration concept: as explained above, we use software patterns to describe proven practices of process-driven SOAs. Patterns have two characteristics which make them useful for model integration across modelling domains:

- Patterns describe recurring solutions in a particular problem domain in an informal and holistic manner. Hence, in contrast to most formal modelling notations, they do not abstract away from details that go beyond a specific modelling domain's abstractions, but instead explain the full solution. That is, if the solution has, for instance, implications for the workflow, the organisation, and the software architecture, all these solution elements are described.
- As proven practice descriptions, patterns encode the recurring themes in the same kinds of models. Hence, they are also a good basis for defining a common meta-model for a modelling domain, because patterns typically describe the established, stable abstractions that are used across different modelling approaches and execution languages.

Because patterns are defined only informally, we use pattern primitives as an intermediary abstraction to represent the primitive concerns in the patterns precisely. At this point, it is very important that we use a common meta-meta-model and a common constraint language to define the meta-models that represent the abstract syntaxes of the DSLs: this way, the primitives can be connected via constraints and also primitives that cut across different models can be defined. The model validator can check all structural properties and constraints in the complete model, even if modelling domains are crossed.

4 Meta-models for SOA integration

There are many modelling domains that play a role for a process-driven SOA. In our tool chain, we have so far concentrated on a subset of these domains that deals with the integration of processes and services. In this context, the following types of languages/models are typically used: component architectures, message flow specifications, workflow or business process languages, programming languages and snippets written in programming languages and business object design models.

For our tool chain, we model both message flow specifications and workflow or business process languages, using extensions of UML2 activity diagrams.²

Component architectures are modelled using UML2 component diagrams. Business object design models are modelled using UML2 class diagrams. In this paper, we will concentrate on examples that illustrate the integration of component architectures and flow abstractions, but the integration with business object design models can be done analogously. Programming language snippets are introduced as individual code (as explained in Section 3.1, cf. Figure 2).

As an example for a meta-model definition let us consider the central flow abstractions: the different models that are relevant for a process-driven SOA come together in various kinds of 'flow' models. There are flow models for long-running business processes, activity steps in long-running processes, short-running technical processes and activity steps in short-running technical processes. Even though these flow models have highly different semantic properties, they share the same basic flow abstraction concept and at the same time they are a kind of glue for all the other models that are involved in a process-driven SOA (such as architecture and design models).

We can define meta-models by instantiating the meta-meta-model classes from Figure 3. In addition to the graphical DSL syntaxes that follow the UML's symbols (see the UML2 standard for details about the UML meta-models and graphical syntax), we also introduce a textual Frag syntax. To get a feel for these textual DSLs, here's a small excerpt of the Frag code specifying some classes and relationships. In the rest of this paper, we omit the textual DSL syntaxes as they are pretty much defined in the same way for the different kinds of meta-models. That is, there is a 1:1 mapping between the textual syntaxes and UML's graphical syntaxes for the various model types.

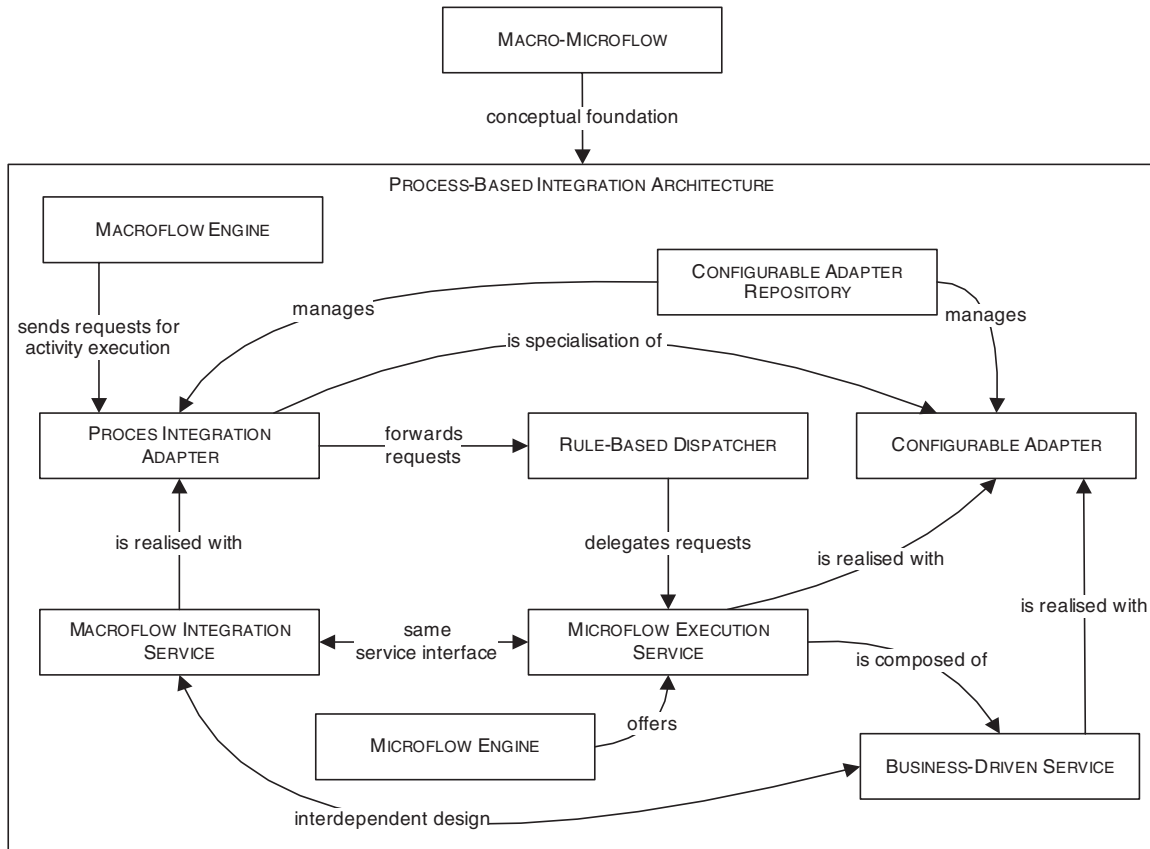
```
MMM::Class create Activity
MMM::Class create ActivityNode
MMM::Composition create ActivityNodes -ends {
  {Activity -roleName activity -multiplicity 0..1
   -navigable 1 -aggregatingEnd 1}
  {ActivityNode -roleName node -multiplicity *
   -navigable 1}
}
MMM::Class create ActivityEdge
MMM::Composition create ActivityEdges -ends {
  {Activity -roleName activity -multiplicity 0..1
   -navigable 1 -aggregatingEnd 1}
  {ActivityEdge -roleName edge -multiplicity *
   -navigable 1}
}
...
```

5 SOA patterns and pattern primitives

We next discuss how to extend the meta-models with pattern primitive extensions. Before we can go into detail, we first give an overview of the pattern language from which we derive the pattern primitives.

5.1 Overview: patterns for process-oriented integration of services

We have mined the pattern primitives from a pattern language for process-oriented integration of services (for details see Hentrich and Zdun, 2006). The patterns and pattern

Figure 4 Overview: pattern language for process-oriented integration of services

relationships are shown in Figure 4. In the pattern language, the pattern MACRO-MICROFLOW lays out the conceptual basis to the overall architecture. The pattern divides the flow models into so-called *macroflows*, which describe the long-running business processes and *microflows*, which describe the short-running technical processes. The PROCESS-BASED INTEGRATION ARCHITECTURE pattern describes how to design a layered architecture, which is following the MACRO-MICROFLOW pattern.

The remaining patterns in the pattern language provide detailed guidelines for the design of a PROCESS-BASED INTEGRATION ARCHITECTURE. The automatic functions required by macroflow activities from external systems are designed and exposed as dedicated MACROFLOW INTEGRATION SERVICES. PROCESS INTEGRATION ADAPTERS connect the specific interface and technology of the process engine to an integrated system. A RULE-BASED DISPATCHER picks up the (macroflow) activity execution requests and dynamically decides based on (business) rules, where and when a (macroflow) activity is executed. A CONFIGURABLE ADAPTER connects to another system in a way that allows to easily maintain the connections, considering that interfaces may change over time. A CONFIGURABLE ADAPTER REPOSITORY manages CONFIGURABLE ADAPTERS as components, such that they can be modified at runtime without affecting the systems sending requests to the adapters. A MICROFLOW EXECUTION SERVICE abstracts the technology specific API of the MICROFLOW ENGINE and encapsulates the functionality of the microflow as a service. A MACROFLOW ENGINE allows for configuring business processes by flexibly orchestrating execution of macroflow

activities and the related business functions. A MICROFLOW ENGINE allows for configuring microflows by flexibly orchestrating execution of microflow activities and the related BUSINESS-DRIVEN SERVICES. To define BUSINESS-DRIVEN SERVICES, high-level business goals are mapped to to-be macroflow business process models that fulfil these goals and more fine grained business goals are mapped to activities within these processes.

Figure 5 shows an exemplary configuration of a PROCESS-BASED INTEGRATION ARCHITECTURE, in which multiple macroflow engines execute the macroflows. Process-integration adapters are used to integrate the macroflows with technical aspects. A dispatching layer enables scalability by dispatching onto a number of MICROFLOW ENGINES. Business application adapters connect to backends.

5.2 Pattern primitives for process and service integration

In this section, we present the pattern primitives for flow abstractions that we have mined from the pattern language explained in the previous section. We will concentrate only on one detailed primitive example and summarise a few other primitives that are needed for the examples in the following sections in Table 1. The other primitives in the full catalogue of primitives have been defined in the same way.

Each primitive is precisely specified in the context of the UML2 meta-model using OCL constraints. To illustrate the precise specification of the primitives let us consider the Macro-Microflow Refinement Primitive.

Figure 5 Example configuration of a process-based integration architecture

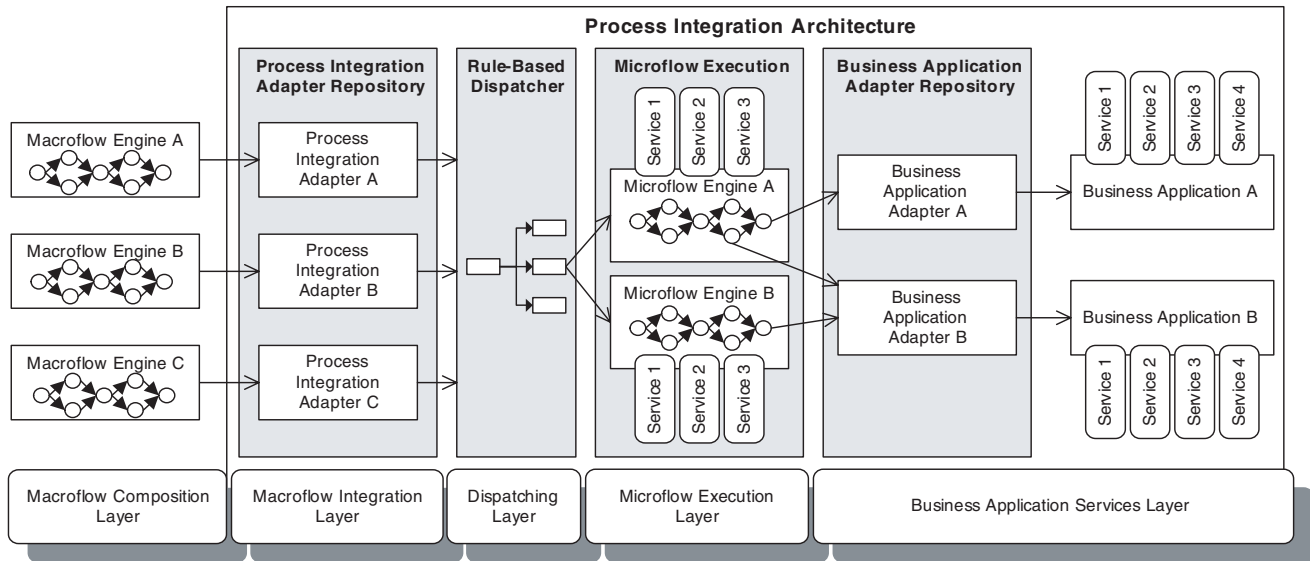


Table 1 Overview: process flow refinement primitives

Primitive name	Description	Modelling solution
Process Flow Refinement	A macroflow or microflow is refined using another process flow	The Activity meta-class is extended with the stereotype Process Flow Refinement, which also introduces tagged values for identifying the refinement
Process Flow Steps	A macroflow or microflow is refined by a number of sequential steps	A specialisation of the Process Flow Refinement stereotype, called Process Flow Steps, is introduced and constrained to be a strictly sequential flow
Macroflow Model	A macroflow can be refined by other macroflows or macroflow steps	Macroflows are modelled by a Process Flow Refinement stereotype, called Macroflow and macroflow steps are modelled as a specialisation of Process Flow Steps, called Macroflow Steps
Microflow Model	A microflow can be refined by other microflows or microflow steps	The microflow model is modelled analogous to the Macroflow Model primitive: The Microflow and Microflow Steps stereotypes are introduced
Macro-Microflow Refinement	Microflow Models are allowed to refine Macroflow Models	The Microflow Model primitive is extended: If refinedNodes of a Microflow is not empty, the Microflow is a refinement of a Microflow, a Macroflow or Macroflow Steps

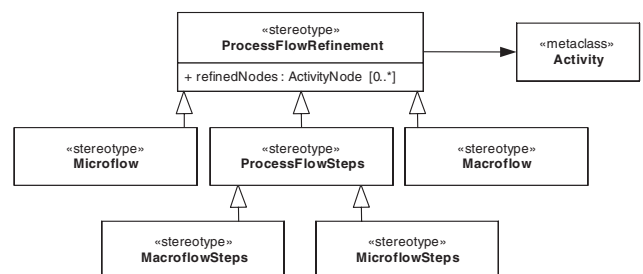
This primitive models the situation that Microflow Models are allowed to refine Macroflow Models. In addition to macroflows and microflows, we must consider the Macroflow Steps and Microflow Steps models, introduced by the Process Flow Steps primitive: A process activity node in a macroflow or microflow can optionally be refined by a number of sequential steps that detail the steps performed to realise the process activity.

To model this primitive, we first must introduce UML2 stereotypes to distinguish the different kinds of refined/refinement activities (see Figure 6) in the UML2 models. We can model the Macro-Microflow Refinement primitive by constraining the Microflow Activities. In particular, if a Microflow activity refines another activity, then this other activity must be itself stereotyped as Macroflow, MacroflowSteps or Microflow. This can be precisely specified using the following OCL constraint:

```
context Microflow inv:
self.refinedNodes->forall(rn |
Macroflow.allInstances()->exists(a |
a.baseActivity = rn.activity) or
MacroflowSteps.allInstances()->exists(a |
a.baseActivity = rn.activity) or
Microflow.allInstances()->exists(a |
a.baseActivity = rn.activity))
```

We have also provided a textual Frag DSL for expressing such constraints. The OCL constraints in UML models can be automatically mapped to this constraint DSL. The task of such constraints is basically to limit the use of the primitives to the acceptable parameters – following the pattern descriptions in which the primitives are used – but no further. For each primitive we have hence described all such constraints (the others are omitted here for space reasons). Thus each primitive describes a precisely specified, parameterisable building block that can be used to specify the corresponding patterns.

Figure 6 Stereotypes for process flow refinements



5.3 Modelling patterns using the pattern primitives for process-driven service integration

The pattern primitives are not yet linked to the patterns. The patterns cannot be themselves precisely specified, but we can identify the pattern primitives that occur in individual patterns. For instance, some of the primitives are mandatory in a pattern, others are optional, still others are only used in specific variants, etc. This mapping of patterns to pattern primitives hence provides us with modelling constructs that can be differently combined for different pattern instances, but must conform to the pattern-to-primitive mapping. If a primitive is used in a pattern instance, all constraints of the primitive must be fulfilled. Hence, the primitives precisely specify the proven practices documented in the patterns. In the rest of this section, we illustrate our approach using the example of modelling the MACRO-MICROFLOW pattern. The other patterns are modelled following the same approach.

The MACRO-MICROFLOW pattern strictly separates the macroflow from the microflow and uses the microflow only for refinements of the macroflow activities. Both in macroflows and microflows we can observe refinements. The different kinds of refinement can be modelled using the Process Flow Refinement primitive. Process Flow Refinement is a generic primitive that can be used for modelling all kinds of process refinements.

Figure 7 shows an example for Macro-Microflow refinement. We describe all models in Figure 7 essentially in the same way. The model integration of the short-running message flow models and the long-running business models is done by extending the models with the respective stereotypes and tag values. After these stereotypes have been defined, the OCL constraints of the primitives enforce that those four models can only be composed in a way that is valid according to the Macro-Microflow Refinement primitive. Figure 7, hence, shows a model conforming to the constraints. Again, the models can be mapped to a textual DSL expressing models defined on top of the DSLs for meta-models, primitives and constraints introduced earlier.

Numerous other variants of the MACRO-MICROFLOW pattern are possible and can be modelled with the primitives introduced. Please note that the flexibility of model assembly through primitives is a very important characteristic of our approach, because it enables us to represent the inherent variability of software pattern solutions.

6 Integrating architectural models

The MACRO-MICROFLOW pattern has implications for short-running message flow models and the long-running business models and we were able to integrate the two model types (and even add macroflow/microflow steps as another kind of modelling abstraction). In our approach, these two kinds of models are modelled with the same model type: activity diagrams. The patterns, however, also have implications for other model types, such as the architectural components in the system or business object models. To model those abstractions, we additionally need to consider architectural abstractions and object-oriented design abstractions. It is typical for patterns that they have implications for different kinds of models; hence the conceptual integration approach, demonstrated in this section for one pattern/primitive example, is also used to model the other patterns and primitives, as well as other modelling domain combinations.

In UML, business objects can be modelled using class diagrams. Architectural abstractions can be modelled via component diagrams. Component diagrams are a specialisation of class diagrams. Therefore, our approach for integrating these models with flow models is very similar for both class diagrams and component diagrams. Hence, we demonstrate our approach only for one of those abstractions in depth: architectural components.

Our approach is to first find suitable pattern primitives for modelling the architectural abstraction corresponding to the patterns. Where further integration between modelling domains is needed, we model the flow model and architectural model with overlapping constraints and primitives that cut across model boundaries.

6.1 Modelling architectural abstractions with pattern primitives

Architectural modelling with pattern primitives follows the same approach as introduced for the flow abstractions. In the context of architectural abstraction, we have introduced similar primitives for architectural patterns (see Zdun and Avgeriou, 2005). Let us consider the Callback primitive as an example: in a process-based integration architecture, different kinds of components are connected. Figure 5 shows an exemplary larger configuration, in which multiple macro-/microflow engines and a dispatcher are used. The

Figure 7 Macro-Microflow modelling example

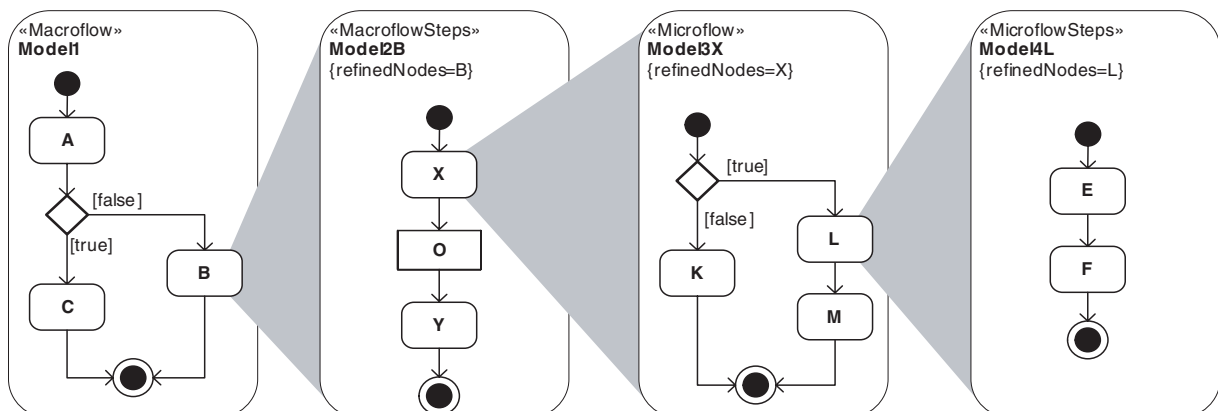
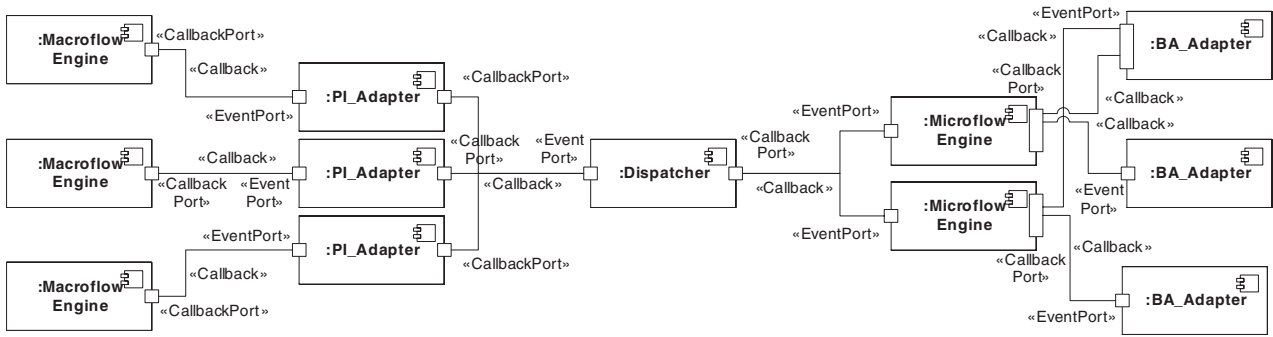


Figure 8 UML2 model for the example configuration



components are typically connected via asynchronous messaging. The modelling support for a process-based integration architecture should allow for flexibly assembling different kinds of architectural models.

The callback primitive (cf. Zdun and Avgeriou, 2005) can be used to model the reactive behaviour in this architecture: a callback denotes an invocation to a component *B* that is stored as an invocation reference in a component *A*. The callback invocation is executed later, upon a specified set of runtime events. Between two components *A* and *B*, a set of callbacks can be defined. To capture the semantics of a callback architecture properly in UML, we propose five stereotypes: **IEvent**, **ICallback**, **EventPort**, **CallbackPort** and **Callback**. Again, we have precisely specified the constraints using OCL (see Zdun and Avgeriou, 2005, for details), and defined the meta-model and constraints in Frag, so that we can use the constraints on components that represent the process-based integration architecture components.

All components in a PROCESS-BASED INTEGRATION ARCHITECTURE are interconnected following the callback style because they use asynchronous communication. The event ports of each layer are listening to events from the higher-level layer, and when an event arrives, they call into the lower-level layer. Once a result is received, it is propagated back into the higher-level layer using a callback. Figure 8 shows an example UML2 model for a callback configuration modelling the situation from Figure 5 in a more precise way. For the architectural primitives we have introduced textual DSLs, just like the ones for the flow abstraction primitives introduced before.

6.2 Integrating architectural and flow abstraction models

In addition to the architectural flexibility of the PROCESS-BASED INTEGRATION ARCHITECTURE pattern, we need to model the pattern’s constraints. If the pattern

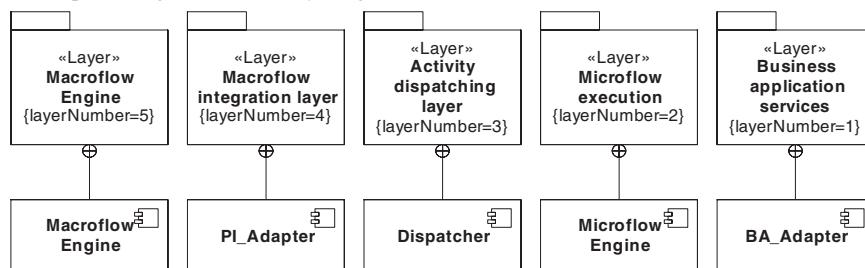
implementation follows the MACRO-MICROFLOW pattern, analogous constraints to the macro-microflow refinement in the flow models must be introduced, such as: components that represent the microflow should not invoke macroflow functionality, macroflow adapters should not be used at the microflow level and vice versa, the dispatcher should only invoke short running microflows, etc.

In this situation, we can use another architectural primitive from Zdun and Avgeriou (2005): layering. Layering describes groups of components and further constrains them. Specifically, it entails that group members from layer *X* may call into layer *X* – 1 and components outside the layers, but not into layer *X* – 2 and below. To model layering in UML2, we introduce the **Layer** stereotype, which specialises the **Group** stereotype (which itself is an extension of the package metaclass). We also impose the following constraints: a component can only be member of one layer and not multiple layers; components who are members of layer *X* may call their fellow components in layer *X*, as well as components in layer *X* – 1 but not in other layers (e.g. *X* – 2 and below). Also, we introduce the tag definition layer number for layers which represents the number of the layer in the ordered structure of layers. Figure 9 shows a UML2 model that extends the model from Figure 8 using the layering primitive.

In this example, we have modelled the integration of the macroflow–microflow refinements at the flow model level with the architectural components by introducing OCL constraints for the refinements between the flow models and by adding similar constraints to the architectural model. That is, the primitive Macro–Microflow Refinement was used at the flow model level and callback/layering were used at the architectural level.

Sometimes this is not enough and it is necessary to extend this strategy and add a direct relationship between the flow models and the architectural models. In such cases, primitives and OCL constraints that contain elements from both model types need to be specified. That is, the OCL constraints cut

Figure 9 Extending the example configuration with layering



across the model types and hence allow for validating that the models do not violate integration concerns between the model types.

7 Related work and evaluation

Our approach supports models for business processes, message flows, OO design and software architecture – and programming language code/snippets provided as individual code in the MDSD tool chain. Our approach is extensible with new model types, especially domain-specific models. We plan to extend our approach in additional relevant modelling domains, such as organisational models or human-interaction models. None of the related approaches offers sufficient support for all these model types. Most related work focuses only on one type of modelling domain.

There are only a few exceptions: Zimmermann et al. (2004) present a generic modelling approach for SOA projects. As in our approach, the approach by Zimmermann et al. is based on project experiences and distills proven practices. The approach also integrates multiple kinds of models for a SOA: object-oriented analysis and design, enterprise architecture and business processes.

Business process management tools, such as Adonis (cf. BOC Europe, 2006) or Aris (cf. Scheer, 2006), describe a holistic model of business process management, ranging from strategic decisions to the design of business processes. They are integrated with standard model types and extensible with new model types. Such tools represent important prior art in the field of model integration. But they do not specifically focus on the field of process-driven SOAs; they are more focused on the business processes. However, an extensible tool suite like Adonis can be used for providing input models for our approach or be extended to model the DSLs.

There are many modelling approaches for business processes, such as Event-Driven Process Chains (EPC) (cf. Keller et al., 1992) and the BPMN (cf. Object Management Group, 2006). Our approach has in common with these approaches that we use the flow abstraction as the central modelling abstraction. Unlike the BPMN, our approach is based on a precisely specified meta-model. Among others, Kindler (2006) has proposed formal semantics for EPCs. In contrast to modelling approaches for business processes, our approach allows to integrate other model types of a process-driven SOA. That is, we take a more ‘integrative’ view than those more specific approaches.

Our approach extends the MDSD concept proposed for instance in Greenland and Short (2004), Stahl and Voelter (2006) with the idea to use one common meta-meta-model for model integration, primitives as modelling constructs based on proven practices and model validation tools for these concepts. In Greenfield and Short (2004, Chapter 13), it is briefly discussed how typical MDSD concepts can be used to support SOA modelling, but only with a focus on web services technology. Essentially, the process description, for example, in BPEL, is seen as a platform for implementing abstractions in a product line and the services are seen as product line assets for systematic reuse. This view does

not contradict our approach, but our approach goes beyond this vision. Through the common meta-meta-model we can integrate any kind of model types; hence, process descriptions are not only a platform, but a first-class model type. In this sense, our approach is also related to the OMG’s MDA proposal, which is a specific MDSD approach focusing on interoperability and platform independence through the distinction of Platform-Independent Models (PIM) and Platform-Specific Models (PSM). Again, our general approach can be used as an extension of MDA, even though our concrete realisation follows more closely the DSL-based MDSD approach (cf. Stahl and Voelter, 2006).

Our approach is not the only approach that is based on proven practices, but only our approach and the workflow patterns approach by van der Aalst et al. (2003) combine proven practices and precisely specified models. The workflow patterns are formalisable constructs (e.g. formalised in the Petri-net-based language YAWL). In YAWL, the workflow patterns are provided as language constructs; hence in the workflow patterns approach the flexibility of assembly of pattern primitives is not (yet) supported, because the variation points offered by the primitives are not offered by the workflow patterns. To support a similar approach as ours, it would be necessary to mine higher-level patterns in workflows that provide guidance on how to assemble the workflow patterns to larger structures.

Some other approaches define particular aspects of service or business process composition using formal specifications, such as the activity-based finite automata-based approach by Gerede et al. (2004) or the interval temporal logic approach by Solanki et al. (2004). Desai et al. (2005) propose to abstract business processes using interaction protocol components which represent an abstract, modular and publishable specification of an interaction among different partner roles in a business process. These approaches aim at model-based verification. Our approach is not designed for this goal. Of course, it is possible to define verifiable models through meta-models and extend our approach, but this has not yet been the focus of our work.

8 Conclusion

In this paper, we have introduced a concept for model-driven development of process-driven SOAs that is based on proven practices. We have especially focused on the aspect of model integration by introducing an approach that is based on a common meta-meta-model from which concrete meta-models for DSLs are derived. In the different DSLs and their respective meta-models, proven practices (described as software patterns) are precisely specified as modelling primitives and their constraints can be validated for all instances of all different meta-models. We have shown in the examples how to integrate message flow models, business process models and architectural models. The approach is, however, applicable for all other kinds of process-driven SOA models for which a precise meta-model is or can be specified. Our tools and DSLs can be flexibly used in model-driven development for precisely specifying process-driven SOAs, validating the models and code generation for executable languages.

References

- BOC Europe (2006) 'Adonis', Available at: <http://www.boc-eu.com/>.
- Desai, N., Mallya, A.U., Chopra, A.K. and Singh, M.P. (2005) 'Interaction protocols as design abstractions for business processes', *IEEE Transactions on Software Engineering*, Vol. 31, No. 12, pp.1015–1027.
- Gerde, C.E., Hull, R., Ibarra, O. and Su, J. (2004) 'Automated composition of e-services: Lookaheads', *Proceedings of the International Conference on Service Oriented Computing (ICSOC 2004)*, New York, NY pp.252–262.
- Greenfield, J. and Short, K. (2004) *Software Factories: Assembling Applications with Patterns, Frameworks, Models and Tools*, J. Wiley and Sons Ltd.
- Hentrich, C. and Zdun, U. (2006) 'Patterns for process-oriented integration in service-oriented architectures', *Proceedings of 11th European Conference on Pattern Languages of Programs (EuroPlop 2006)*, Irsee, Germany.
- IDS Scheer (2006) 'Aris platform', Available at: <http://www.idsscheer.de/germany/products/53956>.
- Keller, G., Nuettgens, M. and Scheer, A-W. (1992) 'Prozessmodellierung auf der grundlage ereignisgesteuerter prozessketten (EPK)', *Technical Report Veroeffentlichungen des Instituts fuer Wirtschaftsinformatik (IWi)*, Heft 89, Universitaet des Saarlandes.
- Kindler, E. (2006) 'On the semantics of EPCs: Resolving the vicious circle', *Data and Knowledge Engineering*, Elsevier, Vol. 56, No. 1, pp.23–40.
- Object Management Group (2006) 'Business process modelling notation (BPMN)', Available at: <http://www.bpmn.org/>.
- Schmidt, D. and Buschmann, F. (2003) 'Patterns, frameworks, and middleware: their synergistic relationships', *25th International Conference on Software Engineering*, pp.694–704.
- Solanki, M., Cau, A. and Zedan, H. (2004) 'Augmenting semantic web service descriptions with compositional specification', 'WWW '04: *Proceedings of the 13th International Conference on World Wide Web*, pp.544–552.
- Stahl, T. and Voelter, M. (2006) *Model-Driven Software Development*, J. Wiley and Sons Ltd.
- van der Aalst, W., ter Hofstede, A., Kiepuszewski, B. and Barros, A. (2003) 'Workflow patterns', *Distributed and Parallel Databases*, Vol. 14, pp.5–51.
- Vinoski, S. (2003) 'Integration with web services', *IEEE Internet Computing*.
- Zdun, U. (2005) 'Frag', Available at: <http://frag.sourceforge.net/>.
- Zdun, U. (2006), 'Tailorable language for behavioral composition and configuration of software components', *Computer Languages, Systems and Structures: An International Journal*, Vol. 32, No. 1, pp.56–82.
- Zdun, U. and Avgeriou, P. (2005) 'Modelling architectural patterns using architectural primitives', *Proceedings of the 20th ACM Conference on Object-Oriented Programming, Systems, Languages and Applications (OOPSLA 2005)*, ACM Press, San Diego, CA, pp.133–146.
- Zdun, U., Hentrich, C. and van der Aalst, W. (2006) 'A survey of patterns for service-oriented architectures', *International Journal of Internet Protocol Technology*, Vol. 1, No. 3, pp.132–143.
- Zimmermann, O., Krogdahl, P. and Gee, C. (2004), 'Elements of service-oriented analysis and design: an interdisciplinary modelling approach for SOA projects', Available at: <http://www-128.ibm.com/developerworks/webservices/library/ws-soad1/>.

Notes

¹Please note that the OMG's MDA proposal is one specific MDSD approach that has some notable differences to our MDSD approach—especially in its focus on interoperability and platform independence.

²Please note that in both cases, long running business processes and short running technical processes, the UML2 activity diagrams must be extended to depict relevant additional information.