



π SOD-M: Building SOC Applications in the Presence of Non-Functional Requirements

Khalid Belhajjame, Valeria de Castro, Javier A Espinosa-Oviedo, Martin A Musicante, Umberto Souza da Costa, Plácido A Souza Neto, Genoveva Vargas-Solar, José-Luis Zechinelli-Martini

► To cite this version:

Khalid Belhajjame, Valeria de Castro, Javier A Espinosa-Oviedo, Martin A Musicante, Umberto Souza da Costa, et al.. π SOD-M: Building SOC Applications in the Presence of Non-Functional Requirements. International Journal of Web and Grid Services, 2018, 14 (4), pp.400-431. 10.1504/IJWGS.2018.095660 . hal-01867030

HAL Id: hal-01867030

<https://hal.science/hal-01867030>

Submitted on 3 Sep 2018

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.

π SOD-M: Building SOC Applications in the Presence of Non-Functional Requirements

Khalid Belhajjame

University Paris-Dauphine, PSL Research University,
CNRS, UMR 7243, Centre Lamsade, 75016 Paris, France
E-mail: kbelhajj@gmail.com

Valeria de Castro

Kybele Research Group, Rey Juan Carlos University, Madrid, Spain
E-mail: valeria.decastro@urjc.es

Javier A. Espinosa-Oviedo

Barcelona Supercomputing Center (BSC), Barcelona, Spain and
French-Mexican Laboratory of Informatics and Automatic Control
(LAFMIA), Puebla, Mexico/ Grenoble, France
E-mail: javiera.espinosa@bsc.es

**Martin A. Musicante and
Umberto Souza da Costa**

Department of Computer Science (DIMAp),
Federal University of Rio Grande do Norte, Natal, Brazil
E-mail: mam@dimap.ufrn.br
E-mail: umberto@dimap.ufrn.br

Plácido A. Souza Neto*

Federal Institute for Education, Science and Technology
of Rio Grande do Norte, Natal, Brazil
E-mail: placido.neto@ifrn.edu.br
*Corresponding author

Genoveva Vargas-Solar

Univ. Grenoble Alpes, CNRS, Institute of Engineering Univ. Grenoble
Alpes (Grenoble INP), LIG, 38000 Grenoble, France and
French-Mexican Laboratory of Informatics and Automatic Control
(LAFMIA), Puebla, Mexico/ Grenoble, France

E-mail: genoveva.vargas@imag.fr

José-Luis Zechinelli-Martini

University of the Americas-Puebla, Cholula, Mexico and
French-Mexican Laboratory of Informatics and Automatic Control
(LAFMIA), Puebla, Mexico/ Grenoble, France
E-mail: joseluis.zechinelli@udlap.mx

Abstract: Specifying non-functional requirements (NFRs) is a complex task, being usually addressed during the latter phases of the software development process. The late inclusion of NFRs during the software development may compromise the quality of the resulting application. This paper presents π SOD-M, a method and associated tools that (i) allow the early specification of NFRs allowing users to them in an abstract way without having to care about low level details; (ii) embraces the MDA philosophy, generating models (code) whenever possible for reducing the specification and programming effort of the business application logic and its associated NFRs. Our solution has been used in the context of an industrial and real case study.

Keywords: MDA, Non-Functional Requirements, Service-based software process.

Reference to this paper should be made as follows: Khalid Belhajjame, Valeria de Castro, Javier A. Espinosa-Oviedo, Martin A. Musicante, Umberto Souza da Costa, Plácido A. Souza Neto, Genoveva Vargas-Solar, José-Luis Zechinelli-Martini (xxxx) ‘ π SOD-M: Building SOC Applications in the Presence of Non-Functional Requirements’, *International Journal of Web and Grid Services*, Vol. x, No. x, pp.xxx–xxx.

Biographical notes: Khalid Belhajjame is a lecturer (Maitre de Conférences) at the University Paris-Dauphine, where he is a member of the LAMSADE. His research interests lie in the areas of information and knowledge management. He has made key contributions to the areas of pay-as-you data integration, e-Science, scientific workflow management, provenance tracking and exploitation, and semantic web services.

Valeria de Castro is associate professor in the Department of Computing Languages and Systems at the Rey Juan Carlos University sited in Madrid, Spain. She received her M.Sc. degree in Information Systems from the National Technological University, Argentine in 2003 and she got her Ph.D. in Computer Science by the Rey Juan Carlos University in 2007. She is member of Kybele Research Group where leads the Service Engineering research line. Her research interests include Services Engineering, Service Oriented Development, Web Engineering and Model Driven Engineering. She has co-authored several publications in national and international conferences and journals. She has participated in several research projects and has lead the main researcher at the Spanish Network of Service Science from 2010 to 2014.

Javier Alfonso Espinosa-Oviedo is a computer scientist specialized in the domains of databases and distributed systems. His research focuses mainly on large-scale data management and processing based on service oriented

architectures. He currently works as postdoctoral researcher at the Barcelona Supercomputing Centre (BSC), in the Extreme Computing Group, and is researcher at the French-Mexican Laboratory of Informatics and Automatic Control (LAFMIA), in the Database Management Group. His research interests include: Stream Processing, Urban Computing, Digital Humanities, Data Visualization. He has participated in several national and European projects (CASES EU-FP7; S2EUNET FP7-IRSES; POLIWEB PEPS CNRS), and has been invited as visiting researcher in different institutions (POLITO, Italy; Udelar, Uruguay; NUA, China; UDLAP, Mexico) where he has lectured on selected topics of data management.

Martin A. Musicante received his B.Sc at ESLAI, Argentina, in 1988 and his M.Sc and PhD in Computer Science at Universidade Federal de Pernambuco, Brazil, in 1990 and 1996, respectively. He is full professor at Universidade Federal do Rio Grande do Norte-UFRN (Natal, Brazil). He is part of the graduate program in Computer Science at UFRN. Martin was associate researcher at LI - Université François Rabelais Tours (2002-2011) and at LIFO - Université d'Orléans, France, since 2008.

Umberto Souza da Costa received the BS degree in Computer Science from Federal University of Rio Grande do Norte in 1998. He obtained the MS degree in Systems and Computation from Federal University of Rio Grande do Norte in 2000. He received the PhD degree in Computer Science from University Federal de Minas Gerais in 2005. He is currently an associate professor in the Department of Informatics and Applied Mathematics (DIMAp) at Federal University of Rio Grande do Norte. His academic interests include Programming Languages, Service-Oriented Computing and Cloud Computing.

Plácido A. Souza Neto received his B.Sc at CEFET-RN, Brazil, in 2005 and his M.Sc and PhD in Computer Science at Universidade Federal de Rio Grande do Norte, Brazil, in 2007 and 2012, respectively. He is an associate professor at Instituto Federal do Rio Grande do Norte-IFRN (Natal, Brazil). He was part of LIG (Grenoble - France) between 2014-2015 for as a post-doc fellow. He is part of the graduate program in Software Analysis and Development at IFRN. Plácido A. Souza Neto's main current research topics are Software Engineering, Data Analysis, Model Driven Development and Web Services.

Genoveva Vargas-Solar (<http://www.vargas-solar.com>) is senior scientist of the French Council of Scientific Research (CNRS) and member of the HADAS group of the Informatics Laboratory of Grenoble, France. She is deputy director the Franco-Mexican Laboratory of Informatics and Automatic Control (LAFMIA). She is regular member of the Mexican Academia of Computing Science. Her research interests in Computer Science concern distributed and heterogeneous databases, reflexive systems and service based database systems. Particularly, she contributes to the construction of service based database management systems. The purpose is to provide methodologies, algorithms and tools for integrating, deploying and executing a service composition for programing data management functions. Data services composition must respect quality of service criteria (e.g., security, reliability, fault tolerance, evolution and dynamic adaptability) and behaviour properties (e.g., transactional execution) adapted to applications requirements. She conducts fundamental and applied research activities for addressing these challenges.

José-Luis Zechinelli-Martini is titular professor of the Department of Computing, Electronics and Mecatronics at the Universidad de las Américas Puebla (UDLAP) since 2002 and he is currently senior researcher at LAFMIA. He is head of the Data and Knowledge Management Group (DBKM) of LAFMIA. In 2001 he obtained his Ph.D. in Computer Science and in 1997 he obtained his Master Degree in Computer Science at University Joseph Fourier, Grenoble, France. In 1993 he obtained his diploma on Computer Systems Engineering at UDLAP. Between 2003 and 2006 he was director of the CENTIA and head of the PhD program on Computer Science. His research project concerns on the integration of Big Data collections on different infrastructures and the specification of spatio-temporal and visualization query languages for retrieving multimedia and multiform data from distributed services.

This paper is a revised and expanded version of a paper entitled “Supporting Non-functional Requirements in Services Software Development Process: An MDD Approach” presented at SOFSEM 2014, Nový Smokovec, High Tatras, Slovakia, January 25-30, 2014 ?.

1 Introduction

In Service-Oriented Computing ?, pre-existing services are combined to build an application business logic. The selection of services is usually guided by the *functional*^a requirements of the application being developed ????. An important challenge of service-oriented development is to ensure the alignment between the functional requirements imposed by the business logic and the functions actually being developed.

Functional properties are not the only aspect to be considered in the software development process. Non-functional properties, such as data privacy, exception handling, atomicity and, data persistence, need to be addressed to fit in the application. Ideally, Non-Functional Requirements (NFR) should be considered in every phase of the software development. Yet, they are partially or rarely methodologically derived from the specification, being usually added once the code has been implemented. In the context of service oriented computing, the late consideration of non-functional requirements in the development process does not fully preserve the compliance and reuse expectations promoted by the service oriented paradigm.

The literature emphasizes the need for methodologies and techniques for service oriented analysis and design ?. Existing approaches note that the convergence of model-driven software development, service orientation and business processes improvement are key for developing quality software ?. Model Driven Development (MDD) for software systems is mainly characterized by the use of models as a product ?. These models are successively refined from abstract specifications into actual computer programs. A recent literature review concludes that current MDD approaches rarely deal with NFRs ?^b.

^aFunctional properties of a computer system are characterized by the effects they produce in the behavior of the system given a specific input.

^bThe study shows that only 48 of the 129 papers considered in their systematic mapping deal with NFR. Moreover, most of them consider just specific aspects of NFR, while only 5 papers (out of 129) propose a more general approach.

Our work introduces π -Service Oriented Development Method (π SOD-M) to support the specification of non-functional aspects of service-oriented applications, taking into account both functional and non-functional requirements, early in the software development process^c. Our proposal follows the MDD guidelines and proposes models, practices and techniques for the development of service-based applications. π SOD-M proposes the use of *models* to specify a software system at different levels of abstraction. Models are organized according to the guidelines of the Model Driven Architecture (MDA) [?]. The goals of π SOD-M are:

- i. To improve the development process by providing an abstract view of the application and helping to ensure the conformance to its specification.
- ii. To reduce the programming effort through the semi-automatic generation of models for the application, to produce concrete implementations from higher-level models.
- iii. To provide a general way of describing non-functional requirements.

The applicability of our proposal was tested with an industrial case study concerning risk assessment for financial companies as implemented by the ORCA System^d. In this work we show the application of our method π SOD-M to develop a service based application called *FlyingPig* that provides risk assessment as a service. The models presented here were generated as a result of interacting with software developers at GCP Global and trying to fulfill all the non-functional requirements present in the application.

This paper is organized as follows. Section ?? summarizes the general principles of existing works for addressing NFP and associating them to service compositions. Sections ?? and ?? introduce respectively the meta-models and transformation rules of π SOD-M. Section ?? describes the *FlyingPig* case study that we develop for validating our method and discusses lessons learned. Finally, Section ?? concludes the paper and gives research perspectives.

2 Related Work

While Functional Requirements establish *what* is computed by an application, Non-Functional Requirements (NFRs) are concerned with *how* the task is performed. NFRs include aspects such as performance, authentication and quality constraints. These requirements are usually specified by conditions, called non-functional properties. Non-functional properties are also referred to as constraints, quality attributes, quality goals, quality of service requirements and non-behavioral requirements ^{???}. Most research on NFRs focus on the evaluation of compliance by the software system as a whole. In service-based applications, NFRs are related to the application itself as well as to its component services.

In ?? non-functional properties of web services are classified according to three perspectives: *service*, *system* and *business levels*. In ? authors use the terms *non-functional attributes*, *composition model entity* and *model entity* to classify different concepts related to NFRs. The notion of non-functional attribute is used to describe NFRs of the abstract process model. In the lower level, the composition is annotated with non-functional attributes.

^cThe letter π of the name of the method comes from *Policy* which is a central notion to represent NFR's and the way they are enforced at runtime. π stands of the pronunciation of the first letter of the word "policy" which is /p/.

^dThe ORCA System is a trademark of GCP Global (www.gcpglobal.com).

D'Ambrogio ? uses the term *quality category* to group similar *quality characteristics*. *Quality dimensions* are used to quantify an individual characteristic. For instance, the quality category *performance* groups characteristics such as *latency* and *throughput*. The development process is based on MDA and the authors also present a WSDL extension for describing the QoS of web services. A catalog of *QoS characteristics* is provided for the web service domain, including properties such as *availability*, *reliability* and *access control*.

Schmeling et al. ? present an approach and a toolkit for specifying and implementing web service compositions with support to several NFRs. Their approach defines abstraction levels, where the terms *Non-Functional Concerns*, *Non-Functional Attributes* and *Non-Functional Actions* are used at each level (in decreasing order of abstraction). A non-functional concern is a general term used to describe NFRs, such as *security*, *reliability* or *transactional behavior*. Each concern is refined into a set of non-functional attributes, where a non-functional attribute represents some behavior, to be refined into non-functional actions. For instance, *encryption* is a non-functional action which provides the implementation of the non-functional attribute *confidentiality*, which is part of the *security* non-functional concern.

Pastrana et al. ? use a traditional design-by-contract approach ?. The authors use the term *contract* to describe non-functional requirements. Contracts may present pre-conditions, post-conditions and invariants. Also, a contract may define *assertions* associated with *quality properties*. Each service may have as many associated *contracts* as needed.

Chollet et al. ? associate (non-functional) *quality properties* to (functional) activities. They present a security meta-model that takes into account web service compositions. In this work the non-functional requirements considered are *authentication*, *integrity* and *confidentiality*. Each NFR is associated with a service activity.

The authors in ?? use formal methods to define a service-based development process that takes NFRs into account. In ?? ontologies are used to define and model NFRs, as in ?? Business Process Modeling is used for system specification, including NFRs. In the method defined in ?, each task and data item of the application can be annotated with functional as well as non-functional attributes (NFAs). Functional and non-functional attributes are independently defined. They are attached to specific tasks later in the development of the application. NFAs for data considers *value* and *range*, whereas NFAs for tasks include *cost*, *time*, *resources* and *expressions*.

The proposal in ? presents steps to select services by taking QoS information into account: (i) identification of relevant QoS information; (ii) identification of basic composition patterns and QoS aggregation rules for these patterns; and (iii) definition of a selection mechanism of services. The authors consider *performance*, *cost*, *reliability* and *availability*.

Karunamurthy et al. ? define NFRs, such as *cost*, *response time*, *availability*, *security*, *reliability* and *reputation*, as *non-function parameters*. The *Non-Functional Specification Language* is proposed as a domain specific language to express *non-function parameters*.

Liu et al. ? use the term *QoS parameter* to describe NFRs such as *cost*, *execution duration*, *accuracy*, *security*, *integrity*, *availability* and *reliability*. In the same way, Tran et al. ? use the term *QoS policies* to classify similar non-functional requirements.

Li et al. ? associate *dimensions* to *QoS parameters* to classify NFRs. For instance, the *time* dimension is associated to the *execution time* and *communication time* parameters; the *spatial* dimension is associated to the *storage capacity* and *message length* parameters; the *reliability* dimension is associated to the *availability* and *reliability* parameters and the

cost dimension is associated to the *service cost* parameter. Rumpel et al. [?] associate *quality requirements* to *quality properties*. Quality requirements are to be specified as constraints.

Ceri et al. [?] use the notions of *policy*, *rule*, *condition* and *action model* to specify NFRs. Agarwal et al. [?] associate *service policies* to services. Each service may also have *properties*, such as *security* and *reliability*. Ovaska et al. [?] use the terms *quality attribute*, *category*, *conceptual layer* and *importance* to organize and classify NFRs. Other authors do not define specific terms to refer to NFRs; they use terms such as *attribute* [???], *property* [?], *factor* [??], *characteristic* [?], *quality level* [?], and *value* [??].

Despite the different notations found in the literature for classifying NFRs, some NFRs are frequently considered, such as *security*, *performance*, *reliability*, *usability*, and *availability*. However, distinct hierarchies and models are proposed for NFRs, according to different perspectives. We have identified a number of approaches [?????] that use MDD (Model Driven Development) for designing and developing applications.

Besides the existence of these works, it is important to remark that authors continue emphasizing the limited attention of the model-driven proposals for service development to NFRs [?]. As Ameller et al. notes in a recently published literature review, there are “*Not many papers that presented a generic approach to deal with NFRs*”. Those authors also remark the lack of industrial use cases [?].

To the best of our knowledge, and based on the related work analysis, there are no proposals that define a service-oriented approach for the whole development process of systems, considering non-functional requirements. Despite the efforts made to support the new technological proposals for the Web, such as web services, these development approaches address their development processes according to traditional software engineering, with emphasis on the functional aspect of the application.

Our method for developing service-oriented applications aims to overcome limitations of existing proposals by dealing with NFRs. Our approach targets the characteristics of 3 areas¹ to propose a method for developing service-oriented applications considering non-functional aspects. The motivation of our method is to target the specification of non-functional aspects in the early stages of development of service-oriented applications. Our approach specifies *policies* for modeling the non-functional requirements of a service. It separates the specification of non-functional requirements from the main functionalities of the application. This is interesting because once the policies for a given application have been defined they can be reused and/or specialized for another application with similar requirements. In order to show these characteristics this paper introduces a validation use case, which corresponds to a service based application provided by the company ORCA. The application is related to the risk assessment industry.

3 Modeling Reliable Service Compositions with π SOD-M

In this section we present π SOD-M, an MDD-based method for building service compositions with NFRs. Our work provides meta-models for modeling functional and non-functional requirements organized in three levels (Figure [??]): CIM (*Computational Independent Models*), PIM (*Platform Independent Models*) and PSM (*Platform Specific Models*). Given high-level models specified at the CIM level, π SOD-M proposes the semi-automatic refinement of these models, for the generation of a set of models at the other levels of abstraction. The refinement process is driven by transformation rules specified

between the meta-models. The purpose of this section is to present the π SOD-M meta-models, from higher to lower level of abstraction, along with a running example. The transformations between models are presented next, in section ??.

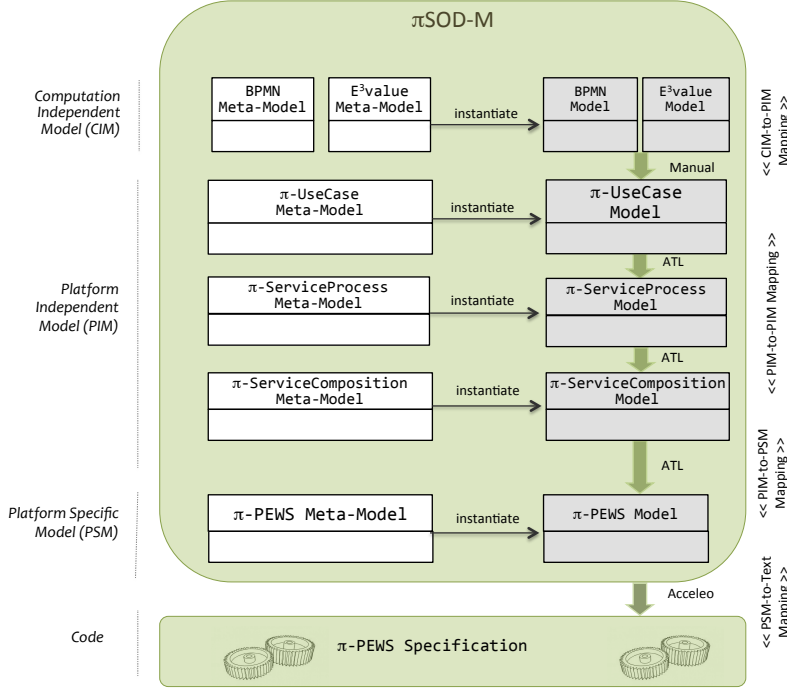


Figure 1 π SOD-M Overview.

3.1 Computation Independent Models

This level focuses on the highest-level view of the system, including its business and requirement specifications. At this stage of the development, the structure and system processing details are still unknown or undetermined. π SOD-M uses the e^3 value ? and BPMN ? meta-models for this purpose.

3.1.1 E^3 value Meta-Model

The e^3 value model identifies the value/information exchange between system components (see Figure ??). This model represents a business case graphically as a set of value exchanges ($\nabla \triangle$) and value activities (rounded boxes)^e performed by business actors (squared boxes). The model is well suited to enhance the understanding of the environment in which an application is being developed. It defines *dependency paths*, showing the value exchange between providers and end users when they ask for a service. A dependency path has a direction and consists of a sequence of linked dependency nodes. It starts with a *start*

^eAccording to ?, a value activity is defined as an activity expected to be profitable to one or more actors.

stimulus node and ends with an *end stimulus* node. Dependency paths may also contain OR and AND elements (both for initiate and join alternative and parallel paths).

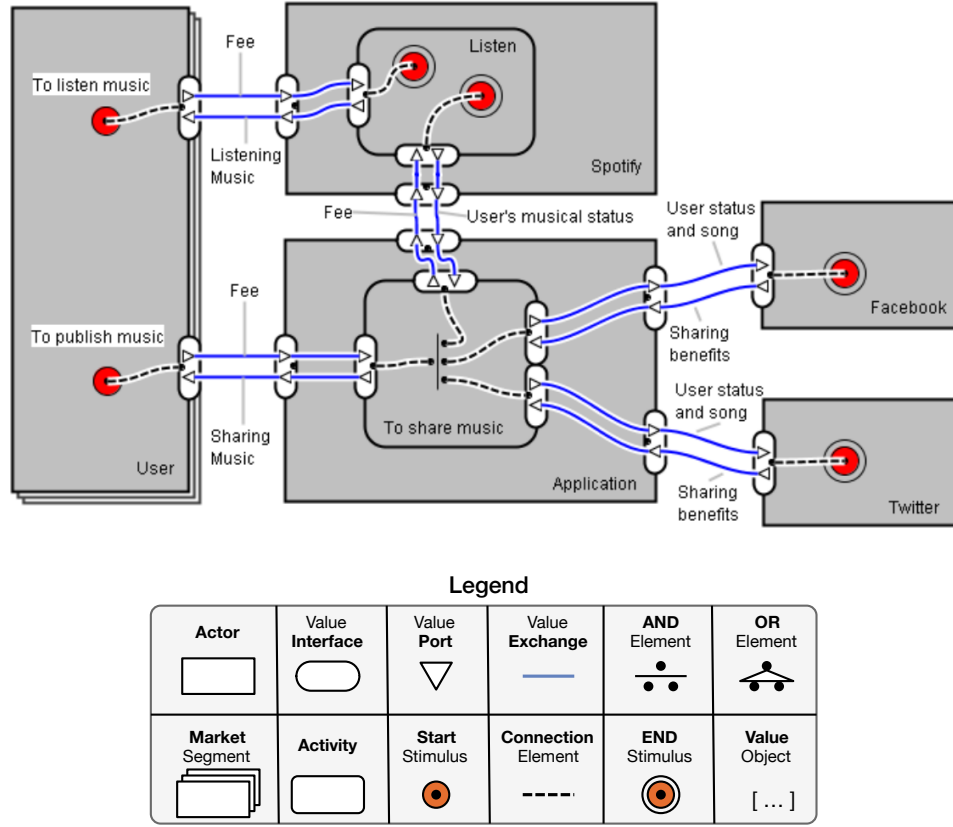


Figure 2 e3value model for “To Publish Music”.

Example 1 (To Publish Music) Let us consider the scenario “To Publish Music”, used as a running example in this section: An organization wants to provide the service based application “To Publish Music” that monitors the music listened by a user during some periods of time and sends the song title to this person’s Twitter and Facebook accounts. In this way, the user will have her status synchronized in Twitter and Facebook (i.e., either the same title is published in both accounts or it is not updated) with the title of the music she is listening in Spotify. The application is based on three external actors (Spotify, Twitter and Facebook). The following (external) services will be used by the application:

- The service Spotify exports a method for obtaining information about the music a given user is listening: `get-Last-Song (userid): String;`
- The services Facebook and Twitter export methods for updating the status of a given user: `update-Status (userid, new-status): String;`

The e^3 value model for “To Publish Music” is shown in Figure ?? The e^3 value model shows Spotify and a private application (which is also a service) that provide services for listening and publishing information about music being listened by users. The private

application interacts with Spotify for obtaining information about the flow of music being listened by a user in return of a fee for a premium subscription. Finally, the private application interacts with Facebook and Twitter for updating the user's status. We can see this interaction as non material benefit sharing (as users subscribe to their networks and are active on them thanks to the private application).

The e^3 value model is used to model the (economic) value exchange among the actors involved in an application. However, the e^3 value model does not help to understand the business process of the application nor the conditions in which the different steps of this process are executed. In order to face this shortcoming, our method proposes the use of the BPMN meta-model as a tool for modeling this aspect of the application.

3.1.2 BPMN Meta-Model

BPMN [†] is a graphical representation that establishes the business process of the application through a high-level workflow^f. The next example illustrates the use of this meta-model.

Example 2 (To Publish Music (cont)) Figure ?? shows the BPMN model of the scenario. It starts by contacting the music service Spotify for retrieving the user's musical status (activity Get Song). Twitter and Facebook services are then contacted in parallel for updating the user's status with the corresponding song title (activities Update Twitter and Update Facebook).




Figure 3 BPMN model for “To Publish Music”.

CIM level models are the basis for the development of the application. The information presented at this level is (manually) refined into PIM-level models of π SOD-M, as described next. Notice that the models are (informally) used to describe both functional and non-functional requirements.

3.2 Platform Independent Models

This level focuses on the system functionality, abstracting the details of any particular platform. The specification defines those parts of the system that do not change from one platform to another. Our method defines three PIM-level meta-models: π -UseCase, π -ServiceProcess and π -ServiceComposition.

[†]Details on BPMN (Business Process Management Notation) can be found in <http://www.bpmn.org/>

3.2.1 π -UseCase Meta-Model

Figure ?? shows the π -UseCase meta-model. This model is used to give the first representation of the application in terms of functionality, and of its non-functional requirements. The notion of *policy* is used to describe NFRs. Figure ?? highlights the concepts related to NFRs. The π -UseCase meta-model extends the UML Use Case meta-model to describe NFRs with the following concepts: BUSINESS SERVICE^g, END CONSUMER, REQUIREMENT, USE CASE, COMPOSITE USE CASE, NON-FUNCTIONAL REQUIREMENT, NON-FUNCTIONAL ATTRIBUTE and CONSTRAINT. An END CONSUMER is represented by an ACTOR. An ACTOR is related to USE CASES (from the original UML definition), as a COMPOSITE USE CASE is a set of actions performed by the system which can be broken into different USE CASES. BUSINESS SERVICE aggregates several USE CASES, a service can be expressed by one or more use cases. The BUSINESS COLLABORATOR concept is represented by PACKAGES. A BUSINESS COLLABORATOR denotes an external service or a system that interacts with the application that is being modelled. Each BUSINESS COLLABORATOR combines the features described in each PACKAGE.

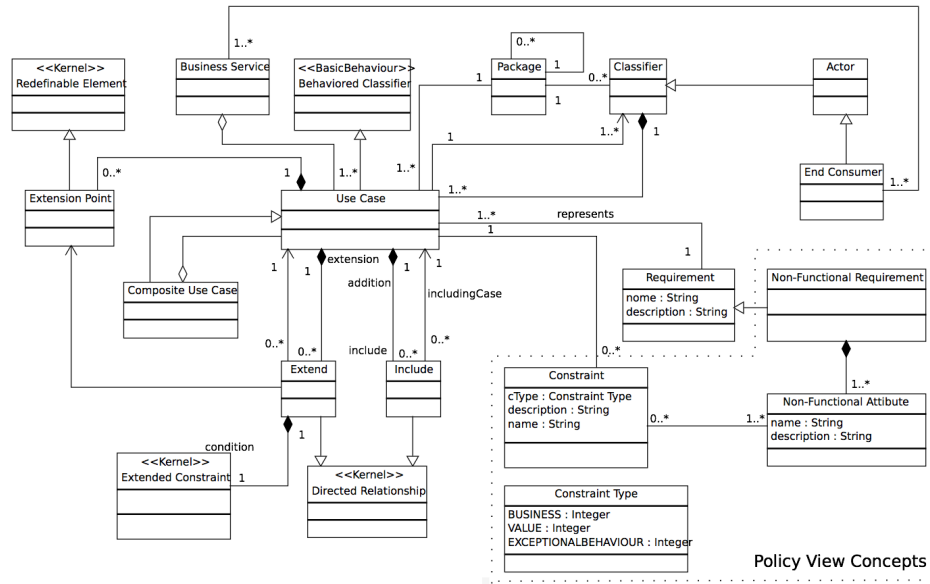


Figure 4 π -UseCase Meta-Model.

The NON-FUNCTIONAL REQUIREMENT and NON-FUNCTIONAL ATTRIBUTE concepts are represented by USE CASES and CONSTRAINTS. A USE CASE may have several CONSTRAINTS. Each CONSTRAINT has a name, description, and a flag to indicate that it has to be dynamically checked. Each CONSTRAINT is represented as a stereotyped (constraint)^h use case.

^gWe use the SMALL CAPS font for referring to classes of a meta-model.

^hWe use the sans serif font for referring to classes defined using a meta-model.

There are three kinds of constraints: (i) A data pair (VALUE CONSTRAINT), represented as the stereotype “value”; (ii) Business rules (BUSINESS CONSTRAINT), represented as the stereotype “business”; and (iii) An EXCEPTIONAL BEHAVIOUR constraint, represented by the stereotype “exceptional_behavior”. The use of these constraints is shown in the next example.

Example 3 (To Publish Music (cont)) Figure ?? shows a π -UseCase model for our example. We consider that, besides the service composition for implementing the application, it is necessary to model other requirements that represent the (i) conditions imposed by service usage – for example, the fact that both Facebook and Twitter require authentication in order to call their methods for updating the wall; (ii) conditions stemming from the business rules of the application logic, (e.g., the fact that the walls in Facebook and Twitter must show the same song title and if this is not possible then none of them is updated).

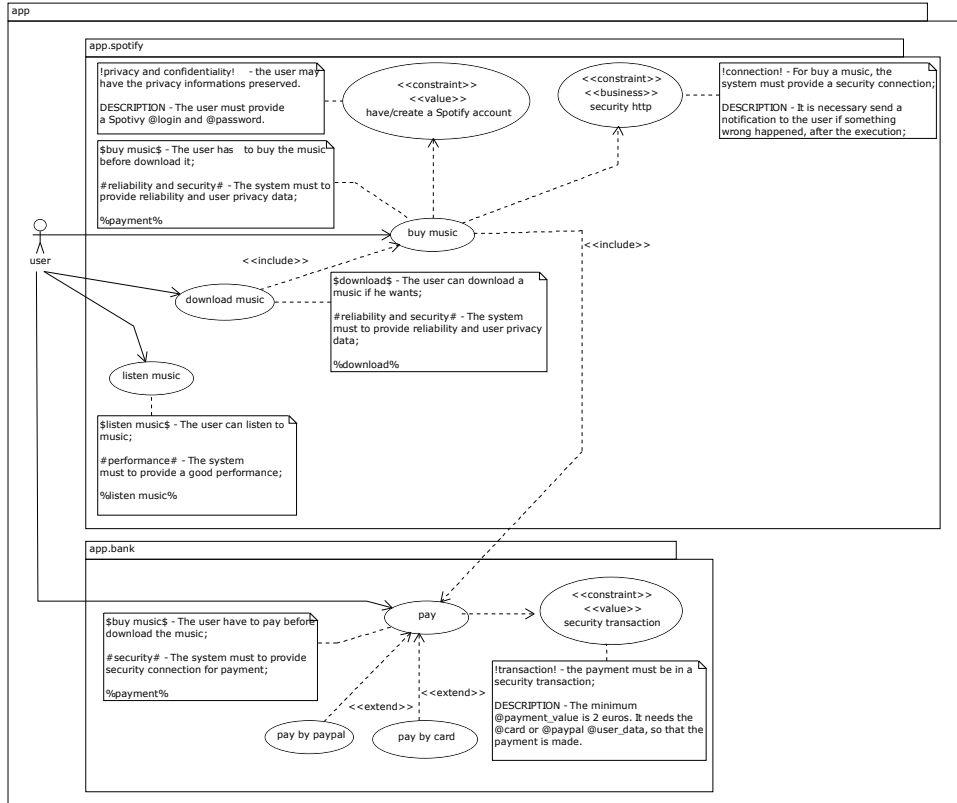


Figure 5 π -UseCase model for “To Publish Music”.

The “To Publish Music” Business Service expects the Facebook or Twitter user status to be changed every time a user starts listening a new song. Therefore, it is necessary to perform a social network authentication with the users data. Each social network uses different services and different forms of authentication. The authentication constraint is required to update a music status. The restriction is stereotyped as a value constraint,

because the user id and password are verified. Figure ?? shows the buy music, download music, listen music and pay use cases. The process of buying a song requires the user private data for a Spotify account, and also a secure connection, represented as value and business constraint, respectively. For payment, the user must provide the data of payment card or PayPal account login and password, represented by a value stereotype. Minimum payment value is 2 euros.

3.2.2 π -ServiceProcess Meta-Model

The π -ServiceProcess meta-model extends the UML activity diagram with the concept of *contract* to represent constraints over data and actions. This concept is used to model groups of constraints in the π -UseCase. In Figure ??, the concepts of the π -ServiceProcess meta-model are: CONTRACT, ASSERTION, EXCEPTIONAL BEHAVIOR, ACTIVITY, SERVICE ACTIVITY, ACTION and CONSTRAINT. The dotted part of the figure defines those concepts related to the representation of NFRs.

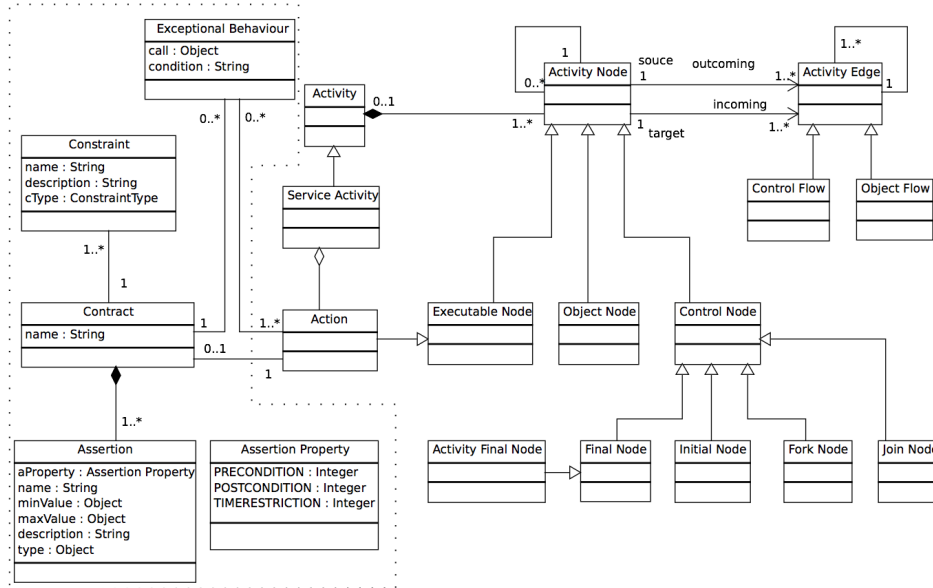


Figure 6 π -Service Process Meta-Model.

This model shows the set of logically related activities to be performed in a service-based application. The activities of this model represent a behavior that can implement the business logic of the application. A π -ServiceProcess model contains three main elements: (i) service process, (ii) service activity and (iii) activity contract. A service activity represents an operation that is part of the execution flow, and it is modeled as an ACTION. An activity contract represents the NON-FUNCTIONAL REQUIREMENT that is also part of the execution flow of a service, identified as a stereotyped activity (assertion). The ASSERTIONS associated to an ACTION compose a CONTRACT. Using the concepts CONTRACT and ASSERTION it is possible to specify each activity service, by defining its pre- and post-conditions. The service that fulfills the CONTRACT and its ASSERTIONS will

be chosen and used in the execution flow. This part of the process requires the verification of some logic conditions. In the present version of the tool, this step requires human intervention.

Example 4 (To Publish Music (cont)) *Considering the example scenario, the contract based process of activities is shown in Figure ???. The buy music and publish music services (update Twitter and Facebook) have pre- and post-conditions assertions that are composed into a contract for each service. The buy music pre-conditions consist in verifying: (i) if the User data are correct; (ii) if the User is already logged in Spotify; (iii) if bank account information is correct and; (iv) if there are enough funds in the bank account to cover the payment. A post-condition ensures the complete transaction and verifies if a notification about the payment authorization was sent both to the user and to Spotify. There are four assertions for the buy music action, and each assertion has been detailed with the assertion property and predicate that must be verified. To update services, depending on each service, there may be different restrictions. As an example, a new verification of user data and message format is appropriate (maximum 140 characters), in the case of Twitter. In the case of Facebook, it is required that the user is already logged in Spotify and these data are the same as Facebook. As post-condition, the application ensures that the Facebook service sends a notification of success. To update Twitter a pre-condition is required, while to update Facebook it is necessary to check a pre-condition and a confirmation notice (modeled as post-condition). As a pre-condition for “twitter update” it is necessary that (i) the music format is correct and (ii) the twitter login and password are correct for the update.*

3.2.3 π -ServiceComposition Meta-Model

In Figure ??, the π -ServiceComposition meta-model provides meta-classes to represent workflows that model business processes. The π -ServiceComposition meta-model extends the UML activity meta-model with the concept of *A-Policy* to group contracts with similar non-functional requirements. For instance, security and privacy restrictions may be grouped into a single policy. This meta-model defines:

- A BUSINESS COLLABORATOR meta-class, to represent the classes of entities that collaborate in business processes by performing some action. An instance of this meta-class is graphically a partition in the activity diagram. A collaborator can be either internal or external to the system. When the collaborator of the business is external to the system, the attribute `IsExternal` of the collaborator is set to **true**.
- ACTIONS, a kind of EXECUTABLENODE, are represented as a class activity instance of the meta-class ACTION. A class action represents some type of transformation or processing. There are two types of actions: (i) a WebService (attribute Type is WS); and (ii) a simple operation called an ACTIVITYOPERATION (attribute Type is AOP).
- The SERVICEACTIVITY meta-class represents classes of composite activity types that must be carried out as part of a business service, composed of executable nodes.
- In order to represent constraint types associated to service compositions, we defined the meta-classes RULE and A-POLICY (see blue meta-classes in the π -ServiceComposition meta-model in Figure ??). We model non-functional constraints by using the notion of *A-policy* ???. An *A-policy* is defined by attributes and rules. The conditions of each rule are evaluated. In case of failure, the actions of the rule will be performed. The meta-class RULE represents event-condition-action rules where the

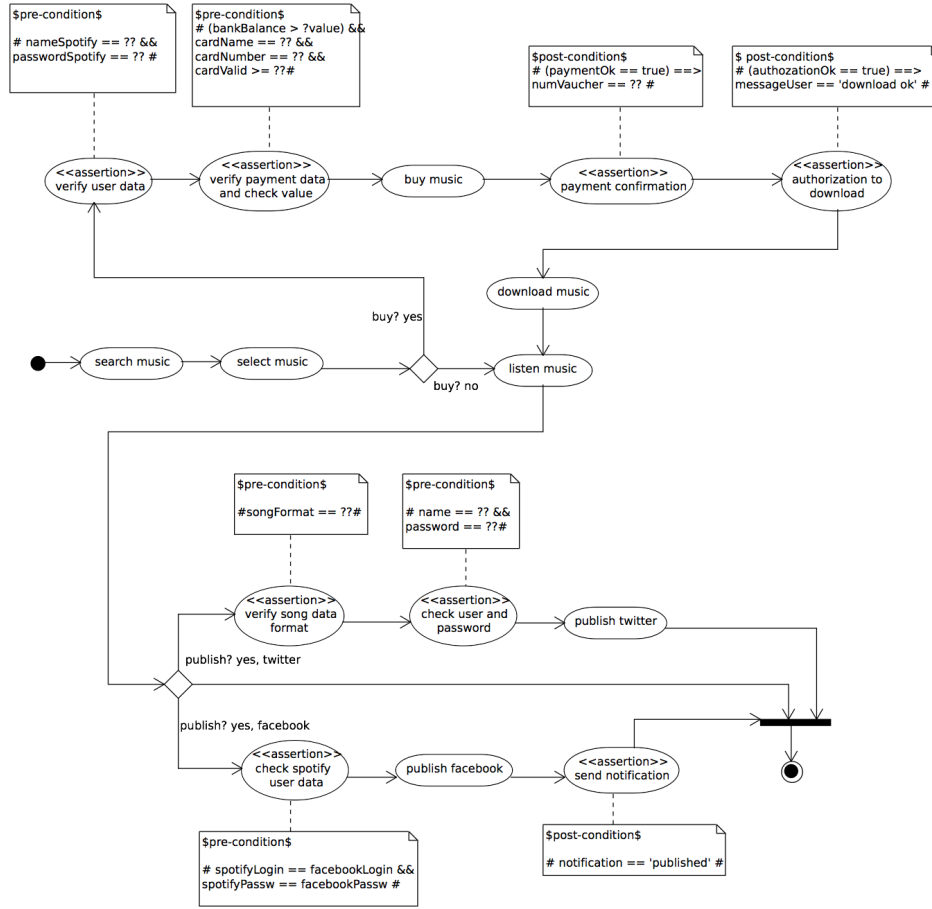


Figure 7 π -ServiceProcess model for “To Publish Music”.

EVENT part represents the moment in which a constraint is evaluated. An *A-policy* defines variables and operations that can be shared by the rules and that can be used for expressing their Event and Condition parts.

Example 5 (To Publish Music (cont)) To illustrate the use of the π -Service Composition meta-model, we define a model for the “To Publish Music” scenario (Figure ??). We model a business process with three service activities: Listen Music, Publish Music and Confirmation. The Publish Music activity calls the Facebook and Twitter services. Both Facebook and Twitter require authentication. Two authentication policies are required, one for Twitter and another for Facebook. In this model, there are three external business collaborators (Spotify, Twitter and Facebook). The model also shows the business process of the application that consists of three service activities: Listen Music, Publish Music and Confirmation. Note that the activity Publish Music calls the actions of two service collaborators namely Facebook and Twitter. Both Facebook and Twitter services require authentication in order to execute methods that read and update the user data. In the example, we associate two authentication policies, one for the open authentication




Figure 8 π -Service Composition Meta-model.

protocol, represented by the class OAuthPolicy at Twitter, associated to the activity UpdateTwitter (see Figure ??). In the same way, the Facebook class HTTPAuthPolicy, for the http authentication protocol is associated to the activity UpdateFacebook. The A-policy OAuthPolicy has a variable Token, used to store the authentication token provided by the service. This variable is imported through the library OAuthPolicy.Token. The A-policy OAuthPolicy defines two rules, both can be triggered by events of type ActivityPrepared: (R₁): If no token has been associated to the variable token, then a token is obtained ; and (R₂): if the token has expired, then it is renewed. Notice that the code in the actions profits from the imported OAuthPolicy.Token for transparently obtaining or renewing a token from a third party. HTTPAuthPolicy implements the HTTP-Auth protocol. The A-policy imports an http protocol library and it has two variables username and password. The event of type ActivityPrepared is the triggering event of the rule R₁. On the notification of an event of that type, a credential is obtained using the username and password.

Once the π -Service Composition Model is defined, it can be transformed into a lower level model (in our case, π -PEWS) to support code generation as described in the next section.

3.3 Platform Specific Models

This level focuses on the functionality, in the context of a particular implementation platform. Models at this level put together the platform-independent view with the specific

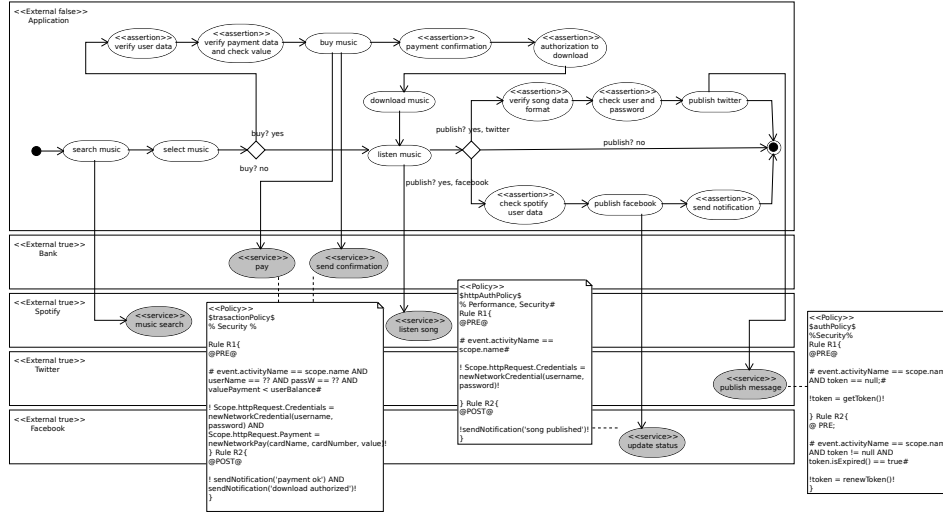


Figure 9 π -ServiceComposition Model for the “To publish music” business service.

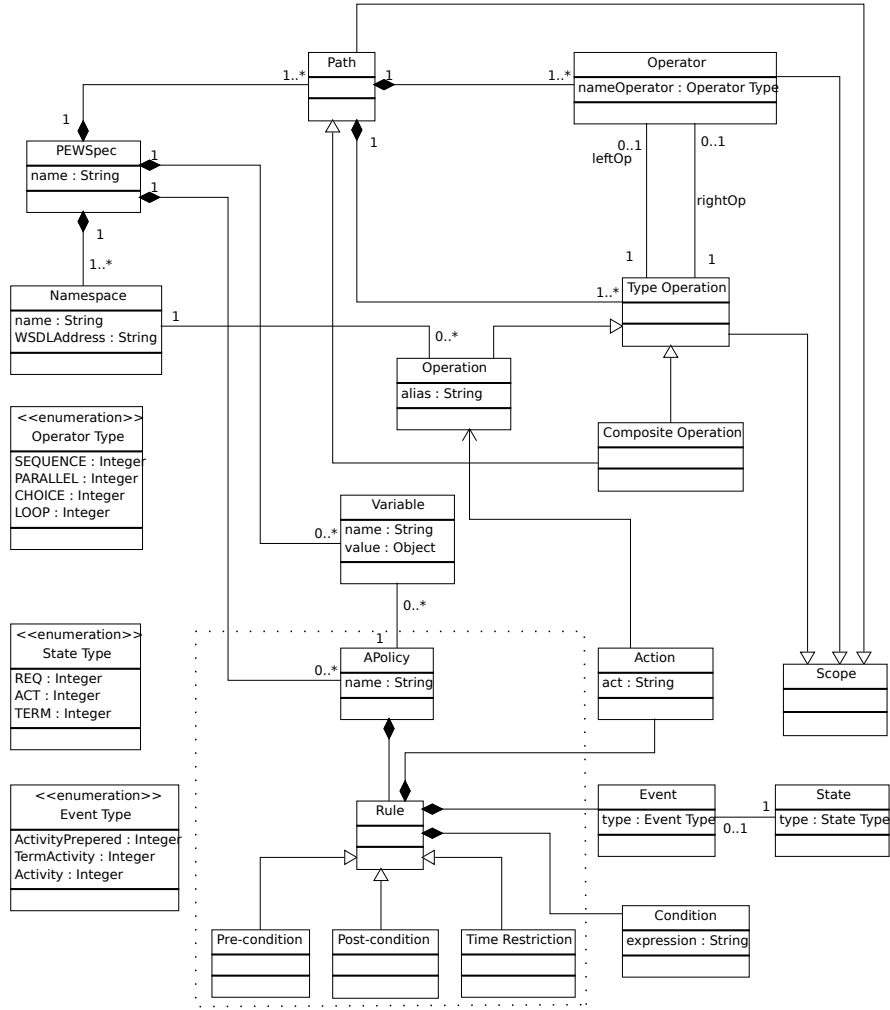
aspects of the platform to implement the system. These models can be automatically translated into actual computer programs. We have defined one meta-model at this level.

π -PEWS provides concepts for modelling service compositions. Instances of this meta-model are textual descriptions of service compositions that can be translated into any service composition language, such as BPEL. PEWS is a notation to express service compositions. The language is based on the notion of Path Expressions and can easily be translated into any actual composition language, such as BPEL. Figure presents the π -PEWS meta-model, where we identify classes to describe:

- Service compositions: NAMESPACE represents the interface exported by a service, OPERATION represents a call to a service method, COMPOSITEOPERATION, OPERATOR and PATH denote service compositions. A PATH can be an OPERATION or a COMPOUND OPERATION. A COMPOUND OPERATION is defined using an OPERATOR. The language defines operators to denote guarded operations ($[C]S$); sequential (.), parallel (\parallel) and alternative (+) compositions; as well as sequential (*) repetition.
- A-Policies that can be associated to service compositions: A-POLICY, RULE, EVENT, CONDITION, ACTION, STATE, and SCOPE.

Figure shows that each A-POLICY is associated to a SCOPE that can be either an OPERATION (e.g., an authentication protocol associated to a method exported by a service), an OPERATOR (e.g., a temporal constraint associated to a sequence of operators) or a PATH. Each A-POLICY groups a set of ECA rules with a classic semantics, i.e, *when an event of type E occurs, if condition C is verified then execute the action A*. In this way, an A-policy represents a set of reactions to be possibly executed when one or several events are notified.

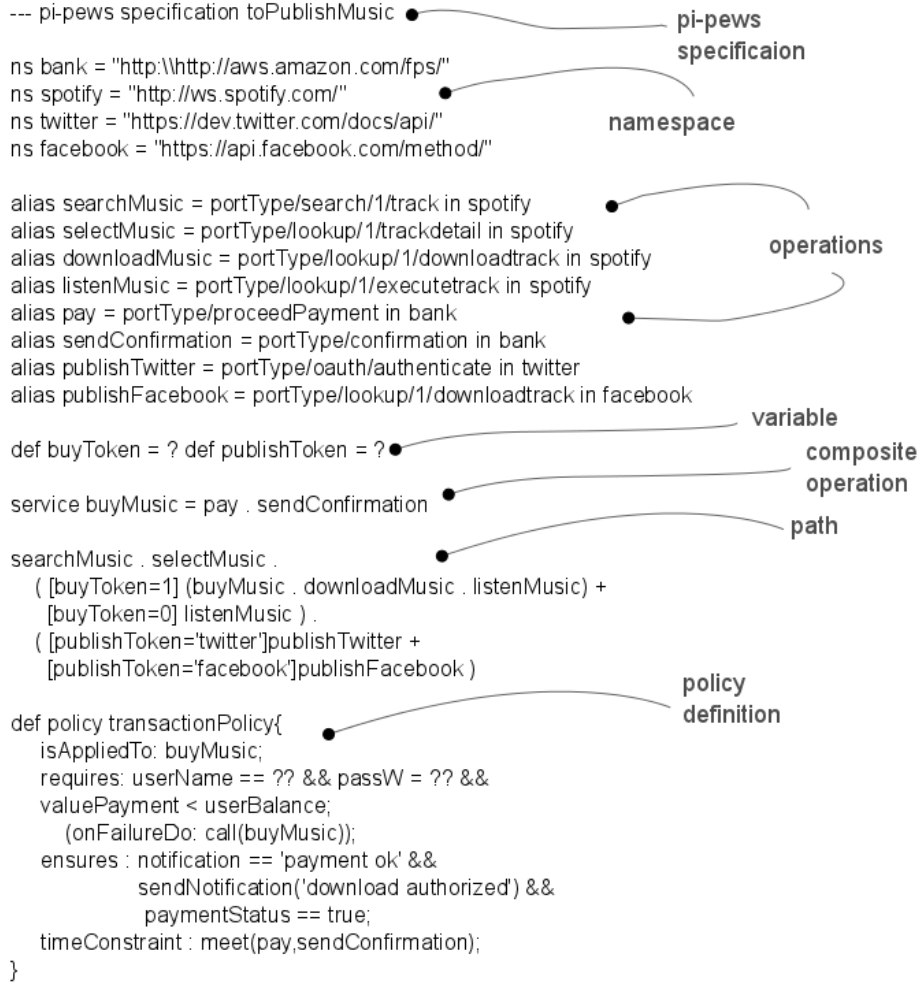
Example 6 (To Publish Music (cont)) Figure shows the π -PEWS code resulting from the π -service composition model of our example. The code contains namespaces and definitions (obtained from the business collaborators of the π -SCM model), a workflow expression (Path) containing operation calls and contracts (derived from the A-Policies).

**Figure 10** π -PEWS Meta-model.

4 Transformation Rules

π SOD-M proposes intra-level transformation rules from π -use case to π -service process models, π -service process to π -service composition models, as well as inter-level transformation rules from π -service composition to π -PEWS models.

These rules were implemented into a semi-automatic tool, the π -SOD-M plug in for Eclipse. At each stage of transformation it is possible to include new constraints. Transformation in the more abstract stages is completely automatic (i.e., from π -UseCase to π -ServiceProcess, and from π -ServiceProcess to π -ServiceComposition). In contrast, the transformation to specific platforms (i.e. from π -ServiceComposition to π -PEWS), the analyst must intervene and check the generation of the specification code because there are programming details that can be interpreted in different manners depending on the target platform. In this sense, the method is more appropriate for those who have background

**Figure 11** π -PEWS Specific and Policy Representation.

knowledge in Information Technology. The transformations from CIM to PIM level models are not automatized, due to the informal nature of the CIM level.

4.1 From π -UseCase to π -ServiceProcess

The refinement of a (composite) π -Use Case model into a π -Service Process model is driven by the principle of expressing a set of π -use cases (i.e., use cases with constraints) in terms of a business process. The resulting model consists of actions related by a control flow and contracts specifying NFPs.

As defined by the π -ServiceProcess meta-model (Figure ??) an ACTIVITY SERVICE consists of a composition of entities of type ACTION. Thus, every π -use case is transformed into an Action of the target π -Service Process model. Every Extend relationship identified in a π -Use Case model is transformed into a Fork node ?. If the Extend relationship concerns just one use case, it is transformed into an Action inside

one flow of the fork node. Otherwise, several use cases are transformed into different Actions that belong to different flows departing from the fork node.

A Constraint associated to a Use Case is transformed into an Assertion. The set of resulting assertions are grouped into a Contract. Constraints are transformed according to their type: BUSINESS constraints and VALUE constraints with the `isExceptionalBehaviour` attribute set to false are transformed into Assertions; VALUE constraints with the `isExceptionalBehaviour` attribute set to true are transformed into Exceptional behaviors.

In order to transform constraints of type Value Constraint, the designer specifies thresholds to be associated to the assertions of a contract. By default, value constraints are transformed into pre-conditions and business constraints are transformed into post-conditions.

The relationships of type EXTEND and INCLUDE determine the way the business process is expressed as a workflow. The generated workflow is composed of Fork and Join nodes, CONTROL FLOW constructors, as well as entities of type ACTION.

Include use case entities are transformed into an Action sequence. A USE CASE element is transformed into an Action. A set of n USE CASES is transformed into an $n - 1$ Object flow elements.

Example 7 (To Publish Music (cont)) *The transformation rules have been applied to the model in Figure ?? to obtain the π -Service Process model for our example (Figure ??).*

The “listen music” use case is transformed into a Service Action that represents a Spotify function to be invoked to play music. For the “publish music” use case, constraints are transformed into a set of assertions that are grouped into a Contract (“publishMusicContract”) associated to the Action “publishMusic”. The use case “download music” includes the payment process to buy the music. Thus, these use cases are transformed into Actions, and a Service Activity that aggregates these Actions. This Service Activity is transformed into a sequence flow on the π -service process model. The same rule is applied to the “publish music” use case, which has two extended use cases, to publish on Twitter and Facebook.

4.2 From π -ServiceProcess to π -ServiceComposition

The transformation of a π -Service Process model into a π -Service Composition model groups Contracts into A-Policies and Actions into Service Activities. Each Assertion of a Contract is transformed into a Rule. Rules concerning the same NFP are grouped into A-Policies. Each Assertion of a Contract is transformed into a Rule:Condition attribute. If the Assertion has a value type, the name and the attributes are transformed into Variables in the target model. The Assertion: `aProperty` attribute can have different transformations, according to: (i) a Precondition is transformed into Pre; (ii) a Post-Condition is transformed into a Post; (iii) a TimeRestriction is transformed into Time.

Some elements of the π -service composition model are obtained from the π -UseCase model: A Package in the π -use case model is transformed into a Business Collaborator; Non-Functional Attributes of the π -use case model are grouped into Non-Functional requirements of the π -ServiceComposition model. These requirements are associated to a Policy (from the π -ServiceProcess model).

Actions and Service activity of a π -Service Process model are transformed into their homonym concepts of the π -Service Composition model. Finally, Actions are grouped into a Business collaborator.

Example 8 (To Publish Music (cont)) Considering the example scenario, the model in Figure ?? was obtained by applying the rules above to the model depicted by Figure ?. The “securityLoginPolicy” consists of a set of Rules that were transformed from the Assertions in π -service process model. The information about Facebook and Spotify (both of them Business Collaborators) come from entity of type PACKAGE in the π -use case model.

4.3 From π -ServiceComposition to π -PEWS

This section describes the PIM to PSM transformations from a π -Service composition model to a π -PEWS model. We distinguish two groups of rules: (i) those transforming service composition entities into workflows; and (ii) those transforming A-Policies into Contracts.

Single actions (Action and Action:name) are transformed into individual service Operations. Complex actions (represented by ServiceActivity and ServiceActivity:name) are transformed into (named) composite operations that define a (named) workflow of the application. Composition patterns expressed using *merge*, *decision*, *fork* and *join* are transformed into their corresponding workflows of the π -PEWS model.

A-Policies defined for the entities of a π -service composition model generate A-policy entities, named according to the names expressed in the source model. The transformation of the rules expressed in a π -service composition is guided by the event types associated to these rules. The variables associated to an A-Policy expressed in a π -service composition model as $\langle \text{Variable:name}, \text{Variable:type} \rangle$ are transformed into Variable, with Name and Type attributes specified from the elements Variable: name and Variable:type of a π -service composition model. Events of type Pre, Post and Time generate, respectively, Pre-conditions, Post-condition and TimeRestrictions.

Example 9 (To Publish Music (cont)) Figure ?? shows the π -PEWS code resulting from the π - service composition model (Figure ??) of our scenario example.

4.4 Implementation

We have developed tools for aiding the user to define and transform the models for all the levels, except for the CIM into PIM level, which should be manually performed. Our tool is implemented as a series of Eclipse plug-ins:

- We used the Eclipse Modeling Framework (EMF)ⁱ to implement the π -Service Composition and π -PEWS meta-models. From these meta-models, we developed plug-ins to support their graphical representation.
- We used ATL^j for implementing the mappings between models.
- We used Acceleo^k to implement the code generation plug-in to generate executable code. It takes a π -PEWS model and generates the code to be executed by the A-Policy based service composition execution environment.

ⁱThe EMF project is a modeling framework and code generation facility for building tools and other applications based on a structured data model.

^j<http://eclipse.org/atl/>. An ATL program is basically a set of rules that define how source model elements are matched and navigated to create and initialize the elements of the target models.

^k<http://www.acceleo.org/pages/home/en>

5 Applying π SOD-M: The *FlyingPig* Case Study

To validate the applicability of our method we have developed a case study concerning risk assessment for financial companies as implemented by the ORCA System¹. Risk assessment is implemented by an interactive business process based on the exchange of questionnaires used to evaluate the risks implied by the client business practice. Examples of business practices are: the conditions and protocols used to perform confidential transactions, the physical security to access reserved areas (such as computing server installations). The information gathered by the questionnaires is used to evaluate whether there are risky practices within the business processes, as well as to propose amendments to these practices. The ultimate goal of the risk assessment is to determine a degree of compliance to existing standards. By analyzing the questionnaires, ORCA detects risky practices, proposes solutions and triggers further assessment processes to ensure that the solutions have been implemented.

Our goal is to model a service based application (*FlyingPig*), for providing risk assessment as a service. In order to provide this functionality, *FlyingPig* benefits from ORCA's legacy services providing storage, assessment and data visualization functions. In the next sections we apply π SOD-M to develop the *FlyingPig* risk assessment system. The models presented were generated as a result of interacting with software developers at GCP Global.

5.1 Computation-Independent Models (CIM)

Figure ?? shows the value model for the *FlyingPig* application. It is a business model that graphically represents a business case as a set of value exchanges (\triangleright and \triangleleft) and value activities (rounded boxes) performed by business actors (squared boxes). We identify two business actors: *ORCA* and *Broker*. Brokers emit requests for risk assessment for one or more companies. ORCA has two value activities which are services that provide an economical benefit: *Identify Amendments* and *Assess Risk Situation*. The values exchanged between ORCA and the brokers are: (i) *Questionnaire and Evidence* filled with information about the client company; (ii) *Amendments* which are ORCA's recommendations, based on the data provided by the answers to questionnaires; (iii) *Evaluation Reports* for the client companies; and (iv) the risk assessment fee. The dependency path in Figure ?? initiates with the need of assessment emitted by a particular company. Once this need has been declared, the value exchanges between ORCA and Broker are triggered. The client company provides ORCA with information (answers to a questionnaire), evidence (to support the information) and a fee (monetary value). In the next step, ORCA suggests amendments (recommendations to change practices) and provides an evaluation report.

Figure ?? shows the BPMN model for the *FlyingPig* scenario. This model is partitioned for better understanding the value exchanging process. The model includes two pools representing the *ORCA* system and the *Brokers*. Brokers have two lanes, the client *Company* and a *User*. The user is a contact member of the company, who coordinates the assessment process. This process involves other members of the company as well. The risk assessment process starts after a request from a company. This corresponds to the value model, in which the start stimulus triggers the whole process. The request leads to the definition of a group of users that will answer questionnaires to evaluate risk. Other tasks

¹The ORCA System is a trademark of GCP Global (www.gcpglobal.com).

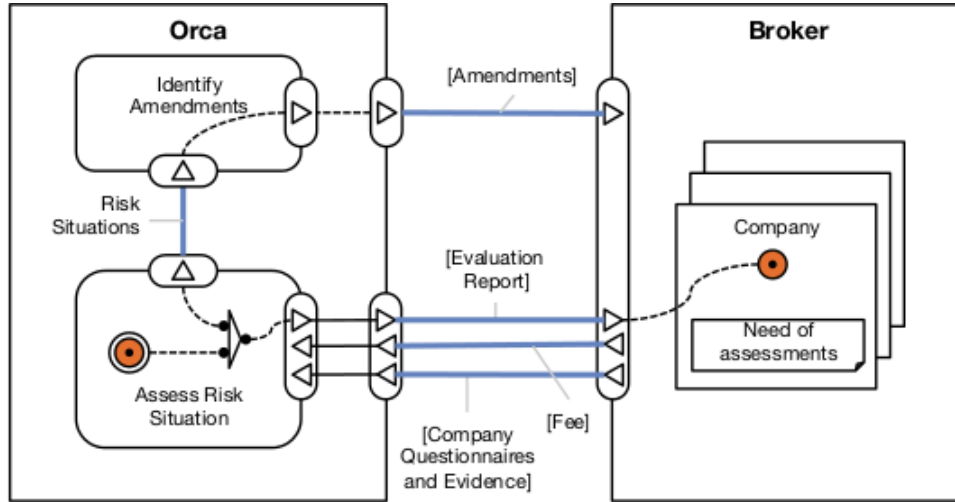


Figure 12 E³ value model for *FlyingPig*.

include amending a “risky situation” as well as producing evidence to show that a specific risk has been eliminated^m.

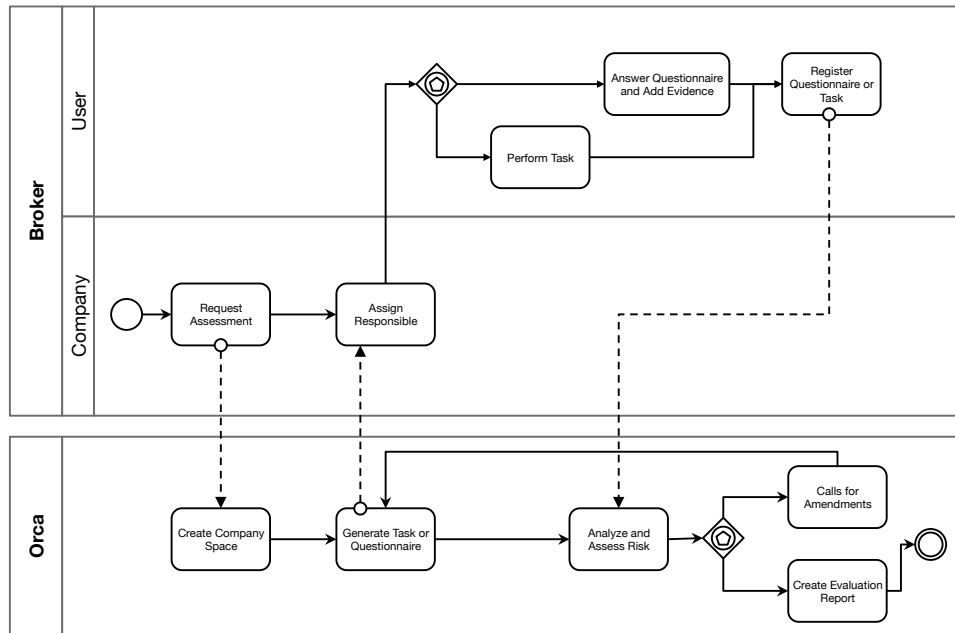


Figure 13 BPMN model for *FlyingPig*.

^mRisky situations include material facts such as not facilitating access to physically disabled people in a bank agency or having an unsecured access to the premises of the company. They can be also immaterial restrictions such as the protocol used for accessing data on the company’s computer server.

Once tasks have been completed, they are stored and analyzed to generate a list of non-compliant situations, associated to their corresponding *calls for amendment* (if needed) or a report specifying a compliance level, incidents and a risk map. During the process of analyzing a questionnaire, the answers to some questions may trigger the generation of additional questionnaires or amendments, that will be scheduled as new tasks. Business processes also have rules and constraints to define their non-functional requirements (NFR):

1. An acknowledgement is due in less than 30 seconds after registering a task or demand for assessment.
2. The system should be able to deal with, at least, 200 users.
3. If the number of requests exceeds 200, *FlyingPig* should implement a load balance strategy for processing the requests.
4. The privileges of the Channel-Broker must be verified *before* the execution of the actions associated to the *designate user in charge* π -use case.
5. The privileges of users must be verified *before* the execution of the actions associated to the *answer questionnaire and add evidence* π -use case.
6. All questionnaires need to be fully answered in order to consider a task as completed.
7. There is a time limit (in days) for amendments required by the system.

5.2 Platform-Independent Models (PIM)

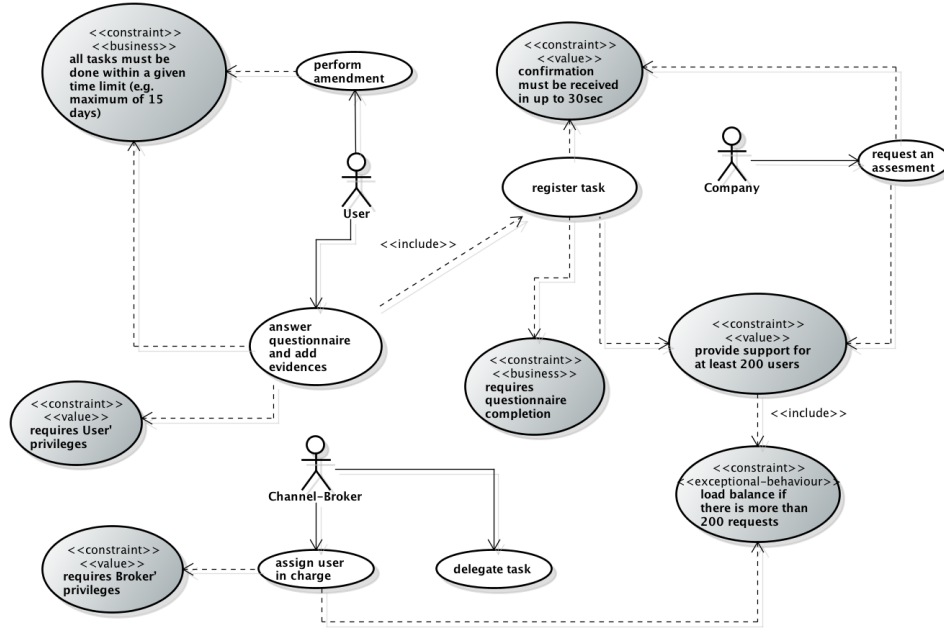
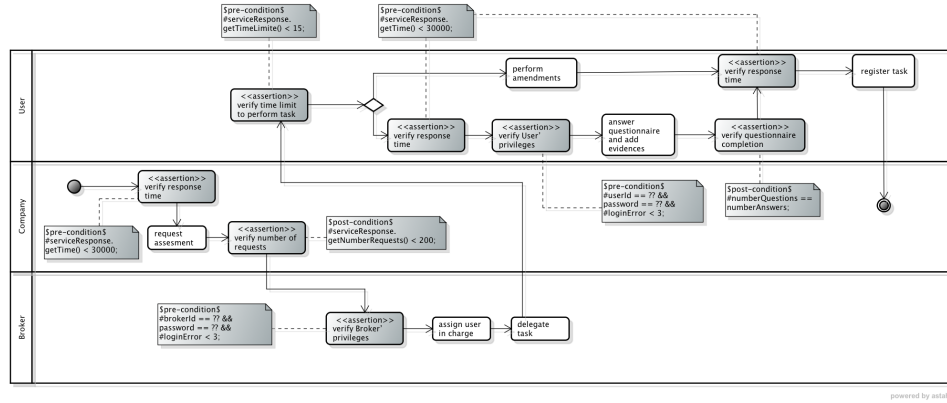
The π SOD-M models of the PIM level for the *FlyingPig* scenario are presented next. These models were manually obtained from the CIM level models.

5.2.1 π -UseCase Model for *FlyingPig*

The π -UseCase model shown in Figure ?? describes the features and constraints of the *FlyingPig* application. In this model, three actors are identified: *Company*, *User* and *Broker* represented as stick figures. *Company* is the actor requesting a risk evaluation as a *Broker* is responsible for coordinating the evaluation process, assigning users to tasks as well as delegating tasks. A *User* is an actor who answers questionnaires according to the current situation of the *Company*. The *User* also produces evidence to support facts and performs the necessary amendments to improve the results of the risk assessment. Each actor is associated to π -UseCases (white ovals in Fig. ??) that describe the main functionality of the system. The π -UseCase model for *FlyingPig* defines six π -UseCases. Each π -UseCase may be associated to (non-functional) constraints (coloured ovals in Fig. ??). Three types of constraints are defined: *value*, *business* or *exceptional behavior*. Constraints are identified by the word <<constraint>> followed by its type.

5.2.2 π -ServiceProcess Model for *FlyingPig*

The π -ServiceProcess model presents the workflow for *FlyingPig* (Figure ??). Actions in this model were obtained by applying the π -use case transformation rules. The *Company*, *Broker-Channel* and *User* actors are transformed into lanes that represent the business collaborators. Use cases are transformed into *actions* and are represented by white boxes. The restrictions associated to π -use cases are transformed into *assertions* (represented by colored boxes) and may be decorated with pre- and post-conditions. We can see that this model refines the concepts defined in the π -UseCase model. The assertions specify those non-functional requirements, as they are seen by the actors. The next step in the development is to add these assertions to the models that specify the *FlyingPig* system.

Figure 14 π -UseCase model for *FlyingPig*.Figure 15 π -ServiceProcess model for *FlyingPig*.

5.2.3 π -ServiceComposition Model for *FlyingPig*

The model in Figure ?? shows the services that provide the functions of *FlyingPig*. The assertions in Figure ?? are implemented as *policies*. These policies express the pre- and post-conditions of the previous model, being associated to the actions of the system.

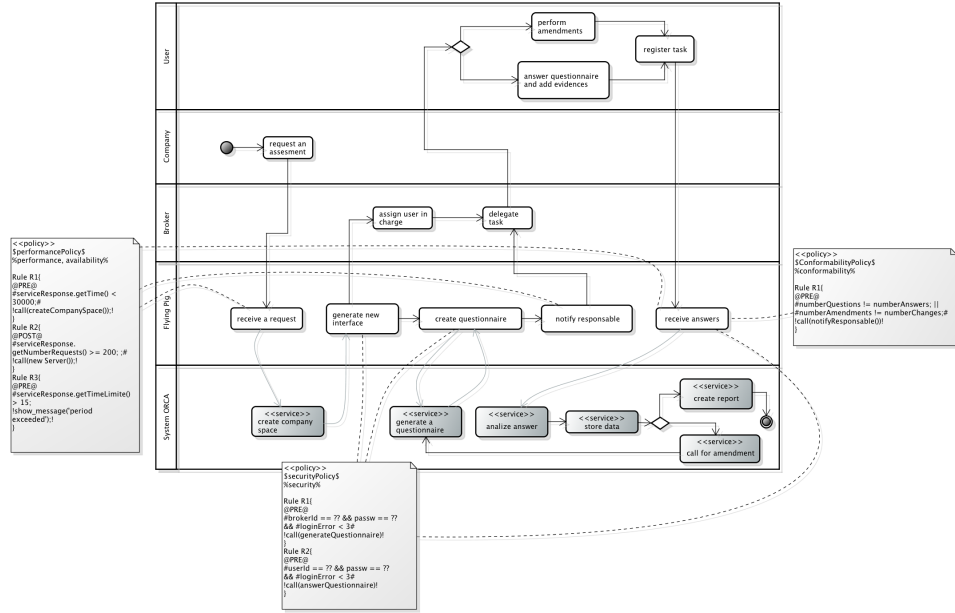


Figure 16 π -ServiceComposition model for *FlyingPig*.

5.2.4 π -PEWS Model for *FlyingPig*

The PSM of our case study is given in Figure ???. This model is obtained by transforming the π -service composition model into a workflow. Note that this workflow is implemented by using BPEL constructors ? and *policy* notions.

```
//Namespaces specify service URI
namespace orca = www.orca.mx/service.wsdl
//Operations
alias createCompanySpace = portType/createCompanySpace in orca
alias generateQuestionnaire = portType/generateQuestionnaire in orca
alias analyzeAnswer = portType/analyzeAnswer in orca
alias storeData = portType/storeData in orca
alias createReport = portType/createReport in orca
alias callForAmendments = portType/callForAmendments in orca
//Services
service receiveRequest(R, Id) = createCompanySpace(R, Id)
service generateNewInterface(Id, NULL) = ...
service createQuestionnaire(Id, Q) = generateQuestionnaire(Id, Q)
service notifyResponsible((Id, Q); NULL) = ...
service receiveAnswers((Id, T); P) =
  analyzeAnswer((Id, T), NULL) . storeData((Id, T), NULL) .
  ((createReport(Id, P) . return(P)) + (callForAmendments(Id,T) . return(NULL)))
//Workflow
receiveRequest(R, Id)
|| (generateNewInterface(Id, NULL)
  . createQuestionnaire(Id, Q) . notifyResponsible((Id, Q); NULL))
|| (receiveAnswers((Id, T); P) . [P != NULL] STOP)*
```

Figure 17 π -PEWS Model for *FlyingPig*.

5.3 Lessons Learned

Through the example we underlined that every application implements functional aspects that describe its application logic. Recall that an application logic refers to routines that perform the activities to reach the application objective. Also there are non-functional properties derived from NFR. They refer to strategies to be considered for the application execution such as security, isolation, adaptability, atomicity, and more. These non-functional properties must be ensured at execution time.

In the context of service-oriented applications, ensuring non-functional properties is challenging due to the nature of their components. Those components are defined by APIs which do not necessarily export information about their internals.

Given a set of services with their exported methods, building service-based applications may consist on expressing an application logic as a service composition. During this task, we must ensure the compliance between the specification and the resulting application. Software engineering methods (e.g., ???) can help to ensure this compliance, particularly when information systems include several sometimes complex business processes calling Web services or legacy applications exported as services.

As WS-* and similar approaches, our work enables the specification and programming of crosscutting aspects (i.e., atomicity, security, exception handling, persistence). In contrast to these approaches, we specify policies for a service composition in an orthogonal way. Besides, these approaches suppose that non-functional properties are implemented according to the knowledge that a programmer has of a specific application requirements but they are not derived in a methodological way, leading to *ad-hoc* solutions that can be difficult to reuse. In our approach, the policies defined for a given application can be reused and/or specialized for another one with the same requirements or that uses services that impose the same constraints.

Wrapping up, the industrial use case presented here provided a real validation context for testing the usability of our method π SOD-M. The software units acting as services provide specific functions and constraints that were combined to build the business logic of the FlyingPig application. This application is representative of the kind of target systems of π SOD-M. The different types of business rules could be modelled as first class citizens and clearly associated to the business logic of the system. Thereby, we could show how non-functional properties can be associated both to the components (i.e., services) and to the global business logic of the system. π SOD-M allowed to derive these non-functional properties thanks to the specialization and transformation of the different meta-models.

We learned and showed that decoupling the description (specification) of non functional properties from the functional aspects, allows developers to focus on the logic of their application, isolating them from the non-functional aspects. Furthermore, it improves modularity by allowing the same application to be run under different non-functional settings, if needed. This can be interesting if business rules change or if constraints associated to services evolve.

Finally, the system FlyingPig had been incrementally built using classic software engineering techniques before applying π SOD-M. It is difficult to compare the initial time to market versus the time to market using our method. The company could not estimate a realistic value of this measure for the first version of the application. We intend to evaluate the quality of experience of programmers, when they will have to build a new version of FlyingPig. The quality of experience related to π SOD-M will be given by measuring (i) time to market when a new module will be designed and implemented in the system (ii)

maintainability of the new module having a separated and clear vision of business logic and non-functional properties. Both two measures are important to the company.

6 Conclusions and Future Work

We presented π SOD-M, a model-driven method for designing and developing reliable service-based applications. π SOD-M extends a previously defined method (called SOD-M) to include Non-Functional Requirements. These requirements are taken into account from the early stages of the software development process. Non-functional constraints are related to business rules associated to the behavior of the application and, in the case of service-based applications, they are also concerned with constraints imposed by the services. Our method includes two CIM-level models, three PIM-level models and one PSM-level model. We implemented the meta-models on the Eclipse platform and we validated the approach by using an industrially inspired use case.

Our case study was developed together with our industrial partner GCP Global to demonstrate the applicability of π SOD-M. The Company is using π SOD-M for the development of their product. The case study presented is a simplified version of their application.

In the future, we expect to: (i) select metrics for assessing the efficacy of π SOD-M; (ii) use the feedback provided by our industrial partner to improve the methodology; (iii) assess the use of π SOD-M in other contexts, in order to investigate the generality of the approach.

Acknowledgements

This research has been partially financed by the Government of the Autonomous Community of Madrid under the SICOMORo-CM project (S2013/ICE-3006), and by the ELASTIC project (TIN2014-52938-C2-1-R), financed by the Spanish Governments Ministry of Science and Innovation. This work is partially supported by INES, grant CNPq/465614/2014-0 (Brazil).