

Architectural Decision Making for Service-Based Platform Integration: A Qualitative Multi-Method Study

Ioanna Lytra¹, Stefan Sobernig², Uwe Zdun¹

¹Faculty of Computer Science
University of Vienna, Austria
Email: *firstname.lastname@univie.ac.at*

²Institute for IS and New Media
WU Vienna, Austria
Email: *stefan.sobernig@wu.ac.at*

Abstract—Nowadays the software architecture of a system is often seen as a set of design decisions providing the rationale for the system design. When designing a software architecture multiple levels of design decisions need to be considered. For example, the service-based integration of heterogeneous platforms and the development of applications on top of those integration services requires high-level as well as technology-, domain-, and application-dependent architectural decisions. In this context, we performed a series of qualitative studies following a multi-method approach. First, we conducted a systematic literature review from which we derived a pattern language for platform integration featuring 40 patterns, as well as a pattern-based architectural decision model. Then, we performed interviews with 9 platform experts from 3 companies for revising the architectural knowledge captured by the pattern language and the decision model. Finally, we participated in a case study and observed the decision-making process to validate the results further. Our observations resulted in 1) a qualitatively validated, pattern-based architectural decision model and 2) a generalized model of the different levels and stages of architectural decision making for service-based platform integration.

I. INTRODUCTION

In recent years, software architecture is less and less seen as only the components and connectors constituting a system’s principal design, and more and more as a set of principal design decisions governing a system [1, 2]. The idea to gather the architectural knowledge about a software system became a focus of the software architecture community. Key in this context is to document not only the components and connectors, but also the rationale for an architectural design using means such as architectural decision models. An architectural decision model (ADM) documents the decision-making process leading to an architectural design in terms of the architectural design decisions (ADDs) made and their relations (e.g., follow-up decisions, implication dependencies). The ADDs and their relations are collected in a structured form, for example, by using text templates and/or diagrammatic modeling based on a meta-model for ADDs [3, 4].

In this paper, we address the decision making for a particular kind of software architectures: software architectures based on one or more software platforms. A *software platform* is a collection of software sub-systems, like communication middleware and databases, and interfaces which together form a reusable infrastructure for developing a set of related software applications. To build a concrete

application by reusing software artifacts in a platform, the platform lays out a customization and configuration process on top of its interfaces [5]. A software platform, therefore, abstracts from details inside and underneath the platform and thereby facilitates developing, maintaining, and deploying domain-specific software applications. Today, many software systems are based on software platforms. Examples include SAP R/3 for enterprise resource planning, the Facebook Platform for social networks, Amazon’s S3 for distributed data storage, Google’s Android and Apple’s iOS platforms for mobile applications, and Mobicents as a reusable VoIP infrastructure.

Architectural design decisions (ADDs) must be captured for each and every architecture when being designed. Documenting ADDs in a disciplined manner is tedious work consuming critical amounts of time. Besides, ADDs depend on the amount of knowledge available at a relatively early time in the software development process; and architects cannot leverage any scaffolding in ADD documentation processes [6]. At the same time, certain design decisions taken in a given technical domain, such as SOA and service-based platform integration, are found to be taken repetitively. This is due to certain design decisions reflecting established design knowledge in the field; or certain technology or development process choices being firm requirements in the field. It has been proposed to base the application-generic knowledge in decision models on software patterns [6, 7] by reusing the recurring knowledge embodied in the pattern descriptions. In this context, we address the following research questions that have, to the best of our knowledge, not been addressed so far:

What are recurring architectural design decisions on service-based platform integration documented by existing software patterns and pattern collections?

Early works on architectural decision modeling (such as [3]) stressed application-specific architectural knowledge, while in more recent works (such as [4]) also decision models for application-generic architectural knowledge have been proposed [8]. When integrating heterogeneous service platforms for supporting domain-specific software applications, decisions ranging from high-level architectural decisions to implementation-technology-dependent and application-dependent decisions must be tackled to specify the system’s architecture. Motivated by these findings on multi-level decision-making, we investigated:

What are the levels of decision making when designing an architecture for service-based platform integration?

To address these two research questions, we conducted a series of three qualitatively-driven studies on architectural decision making. To identify patterns relevant for service-based platform integration solutions, as well as their potential relationships and the relevant forces and consequences of applying the patterns, we performed a systematic literature review [9]. We reviewed in depth 402 patterns from 11 pattern collections. The result was a candidate pattern language containing 29 patterns (plus 11 referenced patterns). From this pattern language we derived a pattern-based architectural decision model. Next, we interviewed platform experts [10, 11] to validate the architectural knowledge documented in the pattern language and the architectural decision model. Finally, we participated in an industry case study [10] on service-based platform integration in the context of a European research project to learn from the industry partners about the decision-making process.

By integrating the results of the three studies to answer the two research questions, we arrived at two major contributions: First, we documented a pattern language and a pattern-based architectural decision model for service-based platform integration [12], validated and refined through expert interviews and a case study. Second, the analysis of the decision-making process has led to a model identifying the levels and stages of architectural decision making in service-based platform integration.

The remainder of this paper is structured as follows. In Section II, we introduce a motivating example for service-based platform integration. Section III describes in detail the steps we followed in our multi-method empirical study, and Section IV presents the results of this qualitative study. In Sections V and VI we discuss the limitations of our approach and the learned lessons from our study respectively. Finally, we discuss the related work in Section VII and summarize our conclusions in Section VIII.

II. MOTIVATING EXAMPLE

To illustrate the research context of service-based platform integration, we give an example (see Figure 1) extracted from the industry case study on industry automation reported in Section IV: Three heterogeneous platforms, a Warehouse Management System (WMS), a Yard Management System (YMS) and a Remote Maintenance System (RMS) are integrated to allow an operator application to utilize the services provided by these platforms. The YMS manages the scheduling and coordination of trucks in a yard, as well as the loading and unloading of the goods from these trucks. The WMS handles the storage of the goods (or storage bins) into racks via conveyor systems. The RMS system is connected to the warehouse to monitor every incident occurring in the warehouse and the yard. It also supports remote communication of operators and workers in the warehouse and the yard.

An operator application uses the services of the three platforms via a domain-specific virtual service platform (VSP)

which performs service-based platform integration. The overall architecture is schematically illustrated in Figure 1. The VSP must handle various integration aspects including interface adaptation between the platforms; integration of service-based and non-service-based solutions; routing, enriching, aggregation, splitting, etc. of messages and events; handling synchronization and concurrency issues, and so on. To design the details of such integration solutions and the applications on top of it, for all the integration aspects high-level architectural decisions as well as application-specific decisions need to be considered. Furthermore, decisions regarding the technologies employed in the platforms, the applications, and the VSP must be made.

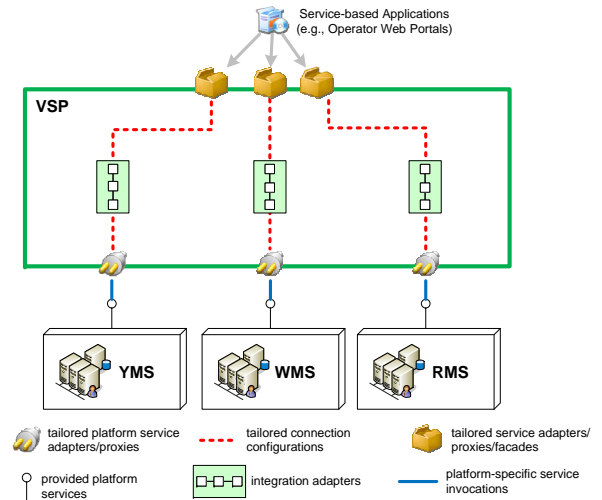


Figure 1. Service-based Platform Integration

III. RESEARCH STUDY DESIGN

When designing our study, we faced the problem of the two research questions being substantially different in terms of their views on architectural decision making; while pattern descriptions relate to architectural decisions and decision drivers content-wise, the issue of decision-making levels requires a process view. A single research method could not accommodate both views. Also, the research questions required both exploratory *and* confirmatory approaches. For example, after having identified candidate pattern material and architectural design decisions as initial results, their *recurring* character should be explained based on practitioners' experience. At the same time, the research questions are closely related, given that different patterns can address different decision-making levels. Consequently, we adopted a multimethod design [13], following a sequential timing for three qualitative study projects.

Figure 2 gives an overview of our research study design: The first method applied was a systematic literature review from which we distilled an initial pattern language and a preliminary architectural decision model according to the procedures in [6, 7]. The systematic review step

was performed by the authors in cooperation with a fourth postdoc researcher (i.e., the authors in [12]; referred to as *reviewers* hereafter). The second instrument was intended (1) to validate and to improve our pattern language and decision model, as well as (2) to explore how decision making in platform integration architectures is performed. For this, we performed semi-structured interviews [11] with experts on three platforms used in industry. The interviews' data were synthesized using the constant comparison method [10] and used to revise the pattern language and the decision model. Based on the integrated interview findings, we designed a case study (e.g., the study proposition) which was then performed as the third and final inquiry step. With this, we were able to document specific design decisions required in the decision-making process. The second and the third research steps were taken by the three authors. In the subsequent sections, we elaborate on the individual study parts in greater detail.

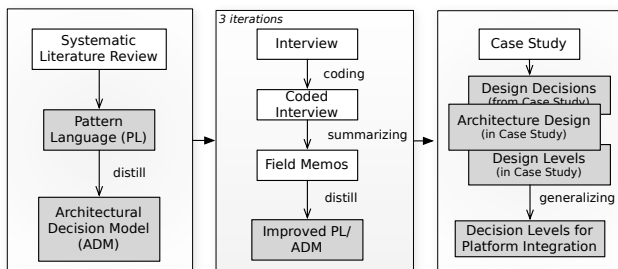


Figure 2. An Exploratory, Sequential, Qualitative Multimethod Design

A. Systematic Literature Review

As an initial step, a systematic review of the pattern literature was conducted to identify, gather, and filter relevant pattern descriptions from pattern sources on distributed system design [14], enterprise application architecture [15], messaging [16], remoting middleware [17], service design [18] and process-driven SOA [19]. In addition, software design [20] and software architecture [21, 22] patterns have been consulted. The goal of this systematic review was to derive a solution space for the technical domain of service-based platform integration, by integrating selected subsets of these pattern collections into a dedicated pattern language. By studying the forces and consequences of the patterns of the resulting pattern language, an architectural decision model was derived with certain patterns entering the decision model as decision alternatives and options (see [6, 7]).

The practice of systematically selecting and integrating parts of existing pattern languages while adding new, context-specific relations between the integrated patterns has been reported before [22, 23]. While pattern catalogues, pattern collections, and pattern languages have been used before to extract and populate decision models [6, 7], our approach based on a systematic pattern literature review (following the guidelines by [9]) is novel. The steps performed in the systematic review process are documented below.

1) *Review Questions*: The questions guiding the systematic review were the following:

RQ_1 Which patterns or pattern collections exist that support designing service-based software system?

RQ_2 Which patterns or pattern collections document integration of service-based software systems?

RQ_3 Which patterns and pattern collections describe the architecture and design of an integration platform?

2) *Pattern Search*: The process for searching relevant pattern descriptions and pattern collections was performed manually by the authors in a coordinated manner. In a first step, the authors held a brainstorming session and identified initial pattern candidates based on their experience in the field. Second, conferences, journals, and book series specific to the pattern community were selected. This selection corresponds to the established stages of publishing pattern material, with pattern descriptions and pattern collections first being disclosed to a pattern audience at a PLoP conference. The PLoP conference series guarantees that the pattern material has been reviewed and edited under the guidance of a shepherd, reviewed by a program committee, and discussed in a writer's workshop. From there, the pattern material may be submitted to a pattern journal, such as LNCS Transactions on Pattern Languages of Programming (TLoP), which subjects the material to the additional scientific review process of an academic, archival journal. Alternatively, pattern material may grow into a book publication. We considered papers originating from 33 conferences proceedings of PLoP and EuroPloP, 2 issues of an archival journal (TLoP), 5 pattern collection books, and 25 pattern books.

3) *Inclusion and Exclusion Criteria*: Patterns and pattern collections, published till February 20th, 2012 were included, provided that the articles met certain minimum requirements: Firstly, the pattern or pattern collection concerns one of the review questions reported before. Secondly, the article presents one or more patterns, in one of the shapes established in the various pattern communities. In particular, the patterns are described in an established pattern form [24] as single patterns, pattern compounds, pattern stories, pattern catalogues, or pattern languages (see [23]). In our in-depth analysis we studied all patterns with regard to the review questions and excluded those that are not addressing one of the review questions. In total we selected 11 sources with 402 patterns for further in-depth analysis¹.

4) *Quality Assessment*: Each pattern publication identified was evaluated according to the following criteria: (1) At least a single version of the pattern collection must have been reviewed by a pattern audience (e.g., a writers' workshop at a PLoP conference, TLoP journal review); (2) At least three known uses are reported; (3) The pattern or pattern collection was considered by the related work [4] on documenting architectural decision in the SOA technical domain.

5) *Pattern Extraction*: The data extracted from the search phase were the publication sources, a categorisation of the pattern material (e.g., single pattern, pattern story, etc.), short

¹Details about the systematic literature review can be found in [25].

Table I
EXCERPT FROM THE INTERVIEW INSTRUMENT

Question	Type
<i>Adaptation and Integration</i>	
1.1 Can the services from the source platform be directly used in the VSP platform?	closed
<i>Interface Design</i>	
2.1 Have the platform services been exposed as services using standard interfaces/technologies?	closed
<i>Communication Style</i>	
3.1 How important is performance for the connection?	open

summarizing descriptions, pattern forces and consequences and other referenced patterns. In accordance with the review guidelines in [9], the extraction was performed by the authors independently from each other.

6) *Synthesis*: We selected 29 out of the 402 patterns in our in-depth analysis for inclusion in the pattern language, plus 11 patterns that are referenced by the pattern language. Detailed results are presented in Section IV-A.

B. Interview Data Collection and Analysis

The interview instrument was carefully designed using the guidelines by Hove and Anda [26]. We performed 3 interviews with 9 experts from 3 different companies that offer the 3 platforms introduced in Section II. The platform experience of the interviewees varied between 1 to 4 years, and most of them had more than 3 years of industry experience. Almost all of them had experience in service technology and were either software designers or developers for the platforms. The interview instrument was mainly based on the decision model and consisted of 4 categories (*Adaptation and Integration*, *Interface Design*, *Communication Style* and *Communication Flow*) and 29 questions. This questionnaire comprised open-ended, as well as closed-ended questions. Out of the 29 questions, 24 questions were based on general architectural knowledge and 5 concerned technical platform details needed for preparing the case study. Table I contains an excerpt from the interview instrument².

The interviews varied between 100 and 150 min. in length. The interviews were recorded on site, using field notes which were then transformed into a structured format through a process of coding. That is, we used a coding scheme to categorize decisions, decision alternatives, and levels/stages of decision making to map the interviewees' answers to concepts such as patterns in our pattern language, relationships between patterns, and correlations between application-generic and application-specific decisions. For example an interviewee's statement that "*the interface can be accessed remotely without any changes*" implies the pattern REMOTE PROXY as a code.

Having analyzed each interview's data, the pattern language and the decision model were reviewed by applying a process of *data saturation* [27]. In particular, new patterns and new connections between the patterns were added to the pattern language. Decision categories and levels/stages

of decision making were documented. Apart from that, patterns and pattern connections that were considered either irrelevant or superfluous were selected as candidates for removal during subsequent iterations.

C. Case Study Research

Our integration case study was both confirmatory and exploratory in nature. Our task was to discuss and to design architectural views together with the platform experts to reflect on the integration architecture based on four integration scenarios. First of all, we studied to which extent the pattern language corresponds to the platform integration domain and the architectural decision model helps navigate the design space. We evaluated the applicability and the appropriateness of the pattern language and of the decision model by making decisions in the context of the case study using these two assets as our main guidance. Apart from that, we analyzed the case study design to gain insights into the decision-making process and to derive hypotheses. The design of a common case study on integrating the three platforms allowed us to define and to explore new decision levels, correlations between application-generic and application-specific design decisions, and different stages of the decision-making process.

IV. RESULTS

A. Pattern Language and Architectural Decision Model

Having synthesized a set of 29 patterns (plus 11 referenced patterns), the authors distilled a pattern language by assigning each pattern to one (or several) of the previously identified thematic categories, resulting in 6 *Adaptation and Integration* patterns, 6 *Interface Design* patterns, 8 *Communication Style* patterns and 9 *Communication Flow* patterns. The full pattern language is reported in a patterns publication [12]. The pattern language documents relations between the patterns, within and between the four categories. In Figure 3, the 6 *Integration and Adaptation* patterns and their relationships are depicted. The patterns are related in four ways: First, a pattern can represent a *variant* of a more generic pattern description (e.g., REMOTE PROXY as a variant of PROXY). Second, patterns can play the role of alternative (exclusive-or) or complementary (inclusive-or) design practices when applied at the same design level (e.g., for integration and adaptation, or denoting communication styles). As an example in Figure 3, the REMOTE PROXY and the INTEGRATION ADAPTER pattern are alternatives for each other, each realizing an integration platform with different capabilities (i.e., direct proxy vs. interface adaptation). Third, adopting one particular pattern (i.e., INVOCATION ADAPTER) *leads to* evaluating another pattern at the same level of abstraction under certain conditions or requirements (e.g., the COMPONENT CONFIGURATOR if runtime adaptability is needed). Finally, some patterns (e.g., PROTOCOL PLUGIN) are used as part of the solution of a higher-level pattern (e.g., REMOTE PROXY) to realize a certain aspect of the solution (i.e., multi-protocol support).

²The complete interview instrument can be found in [25].

In the drafting phase of our architectural decision model, we considered such identified pattern relations as indicators of architectural decisions points to be documented for a single concrete (or even multiple) integration platform architectures. For capturing such recurring, pattern-based architectural design decisions (ADDs), we adopted the following description scheme (see also Figure 3):

- *Decision context:* Arriving at a decision point is motivated by previous decisions taken. Also, the same decision point may be reached several times while constructing an architecture; yet, depending on the given context, different decision options and drivers become relevant. For instance, while Decision 1 in Figure 3 for a given architecture refers to the REMOTE PROXY pattern in the context of the overall integration platform design, the REMOTE PROXY is also relevant in the context of designing the communication flow (e.g., to bridge to a backend platform service).
- *Decision point:* The point of decision making documents the essence of the decision problem. Depending on the decision context (e.g., proxying with or without adaptation), the decision description can refer to the problem and solution statements of decision-related patterns (e.g., REMOTE PROXY and INTEGRATION ADAPTER).
- *Options:* In our pattern-based ADD model, the range of adoptable patterns represent the options space at a given decision point in a decision context. For Decision 1 in Figure 3, applying either the REMOTE PROXY or the INTEGRATION ADAPTER are alternative options.
- *Decision drivers:* The actual drivers for selecting one of the available options can partly be mined from the forces and consequences documented for the decision-related patterns.

Figure 3 illustrates how decisions of our pattern-based ADD model are derived from the pattern language. Table II shows 2 exemplary decisions that have been derived from the excerpt from the pattern language shown in the figure.

B. Interviews and Improvement Iterations

The results from performing the interviews with the platform experts were manifold. The importance of key patterns of our pattern language has been confirmed from the point of view of the independent experts. For example, either PROXIES or ADAPTERS are needed in all cases for invoking the platform services. Some patterns, like SERVICE ABSTRACTION LAYER and EXTENSION INTERFACE were regarded as being useful, but not critical for integrating the platforms at hand. The industry experts agreed that those patterns are rather needed in the field of enterprise applications. In addition, we found that patterns initially omitted in the first version of the pattern language (e.g., PUBLISH-SUBSCRIBER) were used very often by the platform experts. These patterns were added to the pattern language for the next iterations. Apart from that, new connections between the patterns were introduced. Also,

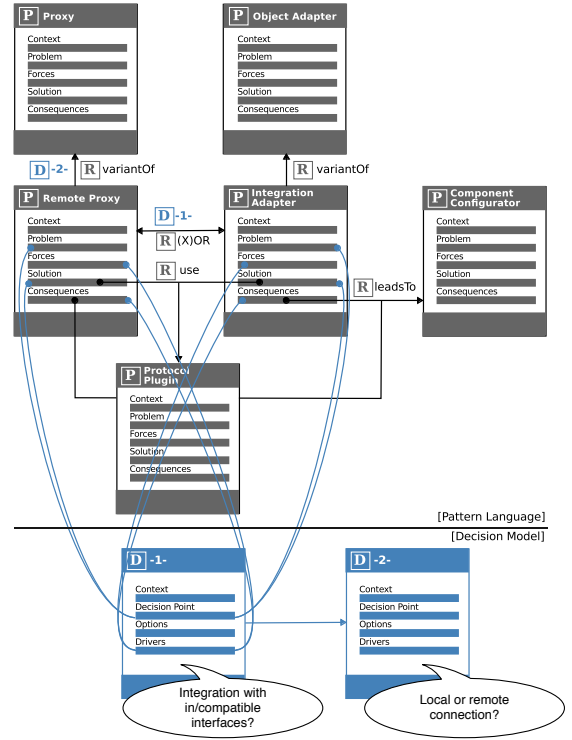


Figure 3. Example excerpt: Pattern Language and 2 Derived ADD Model Fragments

we identified some candidate patterns for removal from the pattern language. For example, the GATEWAY pattern was not considered relevant because its purpose is covered by the SERVICE ABSTRACTION LAYER pattern. Along with the improvements to our pattern language we updated our decision model accordingly.

These interviews were also used as a preparation for the design of the case study. During this process we gathered existing platform services that were combined in four integration scenarios for the needs of an operator application. We noticed that during the interviews the interviewees could not avoid getting into technical details that mainly focused on the technologies used by the platforms and for their integration with other platforms. These technical details introduce constraints that exclude patterns from the pattern language and narrow down the design space. Hence, another important finding of our interviews was that, in the domain of platform integration, decision making has to be performed at multiple application-generic and application-specific levels, and at multiple stages (see Section IV-D).

C. Case Study Design

During the design of the case study, the general architectural decisions were refined during multiple stages; reflecting the single platform technologies, the integration solutions, and the operator application requirements, respectively. The design decisions were documented as instances of the architectural decision model. The outcome of this process was an

Table II
EXEMPLARY DECISIONS IN THE ADD MODEL

Decision Context	Decision Point	Options and Patterns Dependencies
Integration of a platform service	D1 – Which kind of component will be used for integrating the platform service into the service-based integration platform?	<ul style="list-style-type: none"> • None (direct calls from application to platform) • Integration component with same interface (select pattern PROXY or a PROXY variant) • Integration component with a different interface (select pattern ADAPTER or an ADAPTER variant)
An integration component (such as a PROXY or ADAPTER) has been selected for integrating a platform service into the service-based integration platform.	D2 – Is the connection between platform and service-based integration platform a local or a remote connection?	<ul style="list-style-type: none"> • Local (Select local variant of PROXY or ADAPTER, as selected in other decisions) • Remote (Select remote variant of PROXY or ADAPTER, as selected in other decisions)

initial design of the integration solution, for which we used different architectural views to describe the adaptation and the communication levels. The names and the annotations of the design elements signal the use of a specific design pattern in many cases. Due to space limitations, we only present 2 very small excerpts from a component diagram and from a communication flow diagram in Figure 4 to illustrate the different views that have been developed (using the visual notation from [16]). In the next section, we illustrate examples of decision making levels and stages in the context of the case study, which also provide more details about the case study design. ; see <http://www.eaipatterns.com/>

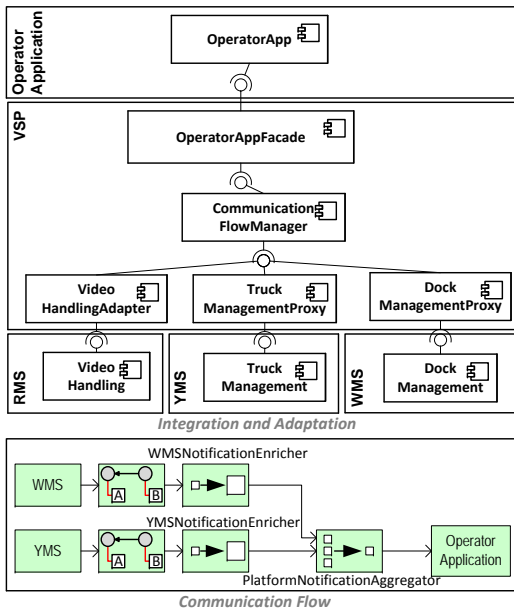


Figure 4. Excerpts from the Case Study Designs

D. Decision-Making Levels and Stages

Approaches to modeling architectural decision making usually focus only on the one or two levels of decision making and knowledge acquisition (i.e., the application-generic and application-specific levels [8]). From analyzing our interviews, we found that in platform-based software development multiple *levels* of decisions and multiple decision times (i.e., *stages*) must be considered. Each decision level is commonly performed by a different stakeholder or

stakeholder group (e.g., platform supplier, system integrator), often working in different organizations. At each stage, the decision space derives from the decisions taken at prior stages. In addition, at each stage the results of decision making at different levels must be taken into account. That is, starting from generic knowledge about a specific form of software platform integration (*Level-1*), architectural and technical knowledge about each of the platforms integrated (*Level-2*), about each of the integration technologies and solutions used (*Level-3*), and about the applications developed based on top of the platforms (*Level-4*) must be considered.

Inspired by the guidelines on multi-stage, multi-level transformations on feature models by Czarnecki et al. [28], we documented the observed decision-making process. Figure 5 depicts an abbreviated example on decisions related to the communication style in our case study. Here, an architectural decision is illustrated using a transformation between two feature models. In our case study, and in the example in Figure 5, the four levels represent decisions to be taken by different stakeholder roles (i.e., integration architect, platform supplier, system integrator, and application engineer). The reusable, pattern-based ADDs form the basis of *Level-1* decisions. The original decision space for each role is modeled as an initial feature model at *Stage-0*. Each subsequent stage sets a discrete decision time in which decisions at different levels of abstractions are addressed at *Level-1* by the integration architect(s) through reviews of historical decisions taken at the other levels. Levels of abstraction are represented by, e.g., appending increasingly specific branches to the decision outcomes (here: the feature trees) for each level.

The flow of decisions in Figure 5 illustrates the process of narrowing down the communication style to be adopted by the service-based integration platform. For instance, at *Stage-1* (in *Level-2*) the architects of the WMS platform supplier have opted for the Windows Communication Foundation (WCF) technology for building asynchronous services (i.e., the `AsyncPattern ChannelFactory` approach). As a result, the integration architects at *Level-1* can refute all the synchronous communication alternatives from the initial architectural decision model. Next, at *Stage-2*, the decisions about the integration middleware (*Level-3*) were considered. Given the predominantly asynchronous communication style adopted by the integrated platforms, the platform integrator chose a message-oriented middleware (ActiveMQ) to drive

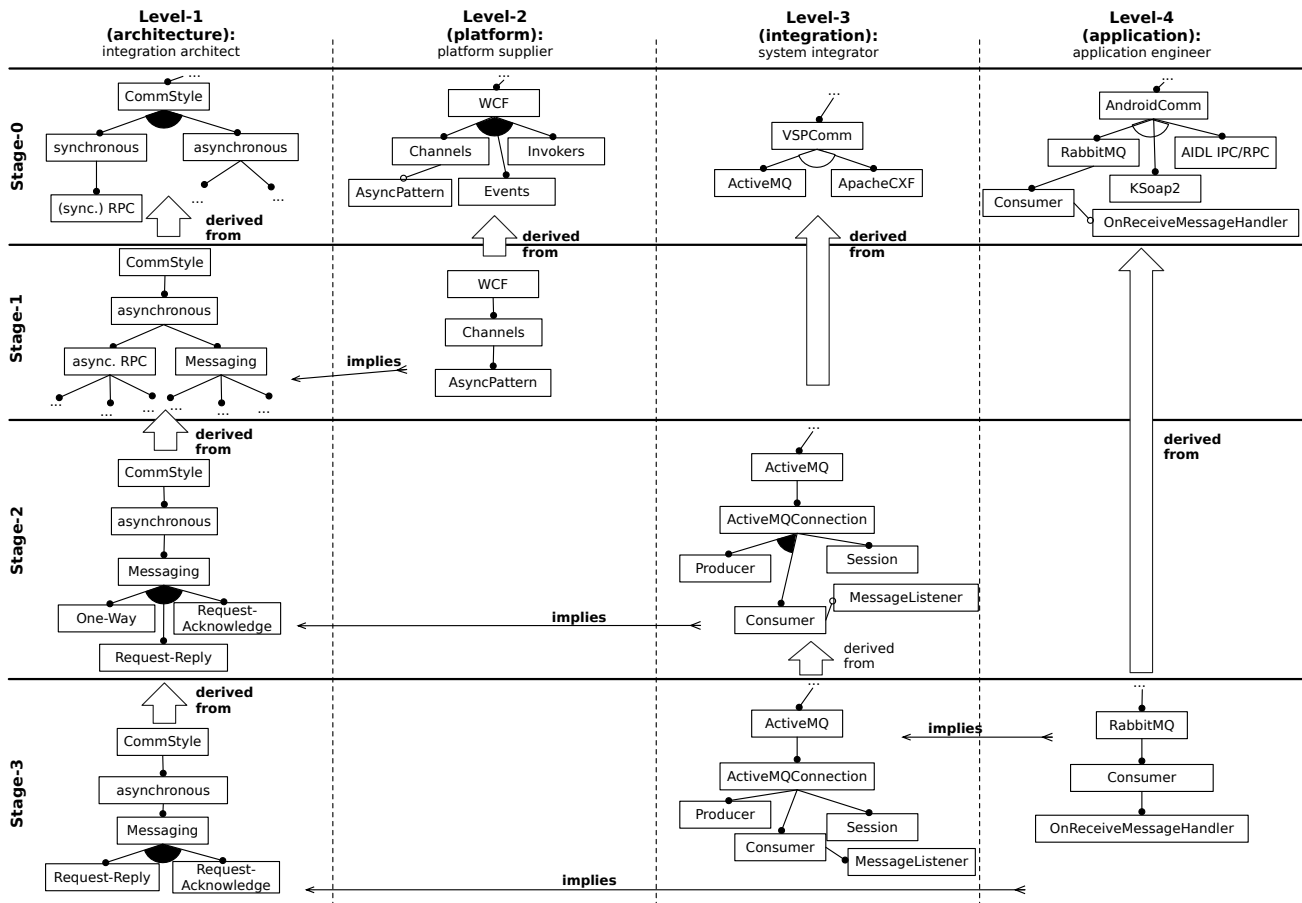


Figure 5. Exemplary Levels and Stages of Decision Making

the integration platform. For the *Level-1* decision space, this meant to refute all the communication style patterns except those related to MESSAGING. Reviewing the decisions taken for the operator application (*Level-4*) at the final stage (*Stage-3*), it turned out that the engineers of the Android-driven application had decided for the RabbitMQ message-oriented middleware and that the application had been built around required notifications about the status of its requests (i.e., using a special callback technique: a `OnReceiveMessageHandler`). Reviewing this decision step meant to specialize the decisions available to the integration architects for the operator application connection even further. That is, any one-way MESSAGING can be excluded for subsequent decision steps. Also, this application-level decision demanded a particular configuration of the message connection in the VSP (*Level-3*), i.e., mandatory message producers and message consumers, to deliver the notifications to the client application.

We have observed this decision scheme in our case study. The dimensions (levels and stages) result from the nature of distributed decision making (e.g., given that the developers of platforms, applications, and integration solutions are often not working in the same organization). However, the

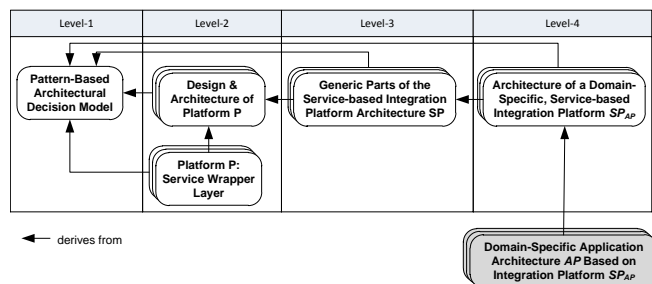


Figure 6. Artifacts and their Architectural Knowledge Derivation Relationships in Different Levels of Decision Making

decisions taken influence the design space left open for later decisions, and each additional stage introduced, excludes or refines decisions that must be made in the subsequent steps. Hence, we propose to extend the existing architectural decision making concepts with support for multiple levels and for multiple stages of decision making along those lines. Our general model of the artifacts at the different decision-making levels and their architectural knowledge derivation relationships are shown in Figure 6.

V. LIMITATIONS AND THREATS

Systematic Literature Review: In conducting the systematic review of pattern literature, we deviated from the original guidelines [9] in important ways: Most importantly, we conducted a purely manual search over a confined selection of conference proceedings, journals, and book series; rather than a semi-automated search process supported by (scientific) search engines. While this practice is consistent with closely related work on pattern-based architecture decision-modeling [7, 8], there is the risk that we have missed relevant pattern material. It simply might be the case that relevant architectural knowledge about service-based platform integration has not yet been documented in pattern form. We consider this threat a minor one because we have found pattern material covering all aspects of the considered cases and we have not found any evidence in our interviews that relevant architectural knowledge is missing. As for the manual search, we explicitly excluded pattern materials which had not been published at any of the pattern community venues.

To include only pattern materials of accepted quality, we relied on the established quality assurance procedures in the pattern community (e.g., shepherding, writers' workshops). To reduce the overall bias of personal judgement (e.g., due to individual research interests, experiences, time constraints), the search step, the quality-assessment step, and the extraction step were performed by each of the reviewers independently from each other. The results were then synthesized in joint re-iterations. In addition, in the overall sequential research design, we validated our resulting pattern language and decision model through interviews with 9 independent industry experts. Despite these efforts, an author's bias cannot be completely excluded.

Interviews: To ensure that the conclusions we took away from the interviews are valid, we discuss how we dealt with the threats to external and to internal validity [11, 29]. As with all qualitative studies, there is the threat to validity with regard to the generalizability of results (external validity) that the small size of considered cases is not enough to generalize the results to other cases of service-based platform integration. We have tried to limit this threat by focusing on patterns as sources of architectural knowledge. In our point of view, it is more likely that similar established best practices are used for other cases of service-based platform integration than if we would have analysed novel, innovative approaches. Apart from that, our interviewees come from three different broad domains (warehouse management, yard management, telecommunications) which of course does not permit us to do any statistical generalization but it offers a broad span of domains in our sample.

As with all qualitative designs, there is the threat to validity that the authors themselves have been an instrument in the study (internal validity) and might have accidentally influenced the study's outcomes or made misinterpretations because of, e.g., limited knowledge, wrong assumptions, and personal bias. We tried to avoid this threat by keeping an

open mind, by encouraging the interview participants to talk and by observing rather than steering the discussion. Also, we included both open-ended and closed-ended questions to get the widest range of feedback. However, it cannot be fully excluded that we influenced the results by our participation in the studies or by our interpretation. Regarding the threats to construct and interpretation validity, we employed *observer triangulation* [11] so that three different researchers were involved as observers and interviewers.

Generalizability: The limitations of the literature search in our systematic-review setup imply that our findings, namely the pattern language and the derived architectural decisions, cannot be reused for similar architecting activities in related SOA settings *with different emphasis*. With the focus on the integration and adaptation view, our review questions and the subsequent pattern search did not cover related SOA concerns such as monitoring or middleware framework design. Pattern material and patterns, even if identified partly, have not been incorporated. Due to the literature filtering and quality assessment, possible connection points to such concerns might have been missed.

The different levels of knowledge among the researchers involved poses another threat. As for the systematic review and pattern extraction, the reviewer group consisted of a pattern expert and senior researcher, two experienced post-doc researchers and a doctoral student focusing on pattern-related research [12]. This heterogeneous distribution of pattern knowledge among the authors certainly affected the preparation of the review protocol, the quality assessment of the pattern material, and the extraction of pattern data for drafting the pattern language. By continuously switching the roles of extractor and checker [30], and multiple iterations over the selected pattern material, this effect might have been substantially lowered, yet not neutralized. The expert opinions collected from the interviews are directly dependent on the 9 experts' experience with the platforms, service-oriented architecting, middleware technologies and software patterns. However, we benefited from a relatively mature expert pool with 7 out of 9 interviewees having more than 3 years of industry experience.

While the aforementioned limitations and risks threaten the generalizability of our results, our methodological approach and the sequential multi-method research design [31] can be applied to comparable research activities on pattern-based architecting.

VI. LESSONS LEARNED

Empirical methods in software engineering are marked by a high level of investigation costs [31]. The time effort required by the authors for preparing and for conducting the systematic review was considerable, caused by face-to-face coordination meetings, numerous phone conferences over a period of two months, the role shifting during pattern extraction etc. Identifying and approaching experienced subjects for the interview and case study steps was facilitated by the involvement of the authors in a large-scale European research project with major industry partners.

In our study, we have learned that using software patterns facilitates iterative architectural decision making [6]. The actual forces and consequences in pattern descriptions document abstracted choices at a given decision point. While our pattern language does not document ADDs for a concrete integration platform architecture, it identifies observed designs for a family of these architectures. In the case study, the pattern language was reused by the architects for sketching a concrete architecture by referencing pattern descriptions from within ADD descriptions. In addition, the pattern form complemented the qualitative research methods used in our study. The structured presentation of patterns and pattern collections facilitated the systematic review. Patterns proved to be an important communication vehicle between the interviewers and the interviewees to bridge the gap between their different technology backgrounds.

At the same time, when preparing our study, we became aware of documented limitations of software patterns as empirical research objects. Especially forcing subjects in experiments into using and applying architectural styles and design patterns resulted in observations indicating an *increased* level of development and maintenance effort. The complexity of learning and comprehending pattern material affected the observations. Such effects have more systematically been reported for research designs involving patterns and artificial software artifacts of lower complexity (toy applications). While the earlier studies adopted experimental designs, we learned two lessons from these. First, our research design should not impose design decisions (in terms of patterns) onto our subjects [31]. Second, the architecting process must be observed in the context of a real, not artificially constructed development project. The first risk was mitigated by confronting the experts not directly with the selected SOA patterns, but rather with design decisions derived from our pattern language. With this, patterns played only an indirect role for the observed architecting process. As for the second risk, the architecture designed in our case study applies to a large-scale software prototype which is to be delivered in the context of a European research project; as such, the architecture is not specific to or was not created for the purpose of our study alone.

VII. RELATED WORK

In this work we create reusable ADDs in the context of platform integration based on design patterns. Many ADD approaches propose prescriptive ADD meta-models [3, 32, 33] that introduce relationships between decision alternatives and activities related to them. By reusing and linking ADDs to patterns as design artifacts [7, 8], documenting architecture decisions and design rationale can be substantially facilitated. For instance, Capilla et al. [33] consider architectural patterns as concrete decision alternatives using a structured Wiki as an ADD documentation tool and a SOA-centric industry case study. Patterns on process-driven SOA [19], on enterprise application integration [15] and on remoting middleware and messaging systems [16, 17] have been documented in the literature, but not focusing

on service-based platform integration as in our pattern language.

In the design decision literature, many approaches consider multiple levels of decision making and distinguish between application-generic and application-specific knowledge. Examples of generic knowledge are software patterns [20, 21], architectural styles, and reusable architectural decisions [7]. Decision models [2, 3] and models created in architecture description languages (ADLs) are examples of application-specific knowledge. A number of approaches (e.g., [3, 7]) capture technology-specific knowledge using decision meta-models and/or decision templates. Zimmermann et al. [32] introduce four levels of an decision-making process: the executive decision level (process and requirement analysis patterns), the conceptual decision level (high-level architectural patterns), the technological decision level (design and remoting patterns) and the implementation decision level (concrete technology options). So far, however, the integration of multiple levels of application-generic and application-specific architectural decisions has not been studied systematically. In our work, we study the inter-decision connections and observe the decision making process in the context of service-based platform integration and we abstract the different levels and stages of this process.

Our multi-level and staged decision making approach is inspired by the staged configuration of feature models [28] which is used for product line configurations. In contrast to this approach, we support the decision making on the architectural level rather than on the concrete software characteristics. In addition, the required knowledge is not completely known during the early stages, but the decision models get concretized and decision models from previous stages get refined through iteration cycles. To the best of our knowledge, an approach to defining stages in architectural decision making does not exist in the literature so far.

VIII. CONCLUSIONS AND FUTURE WORK

In this empirical work, we employed multiple qualitative methods to explore the practice of architectural decision making in service-based platform integration. First, we performed a systematic literature review from which we distilled a pattern language and a pattern-based architectural decision model. In the next stage, we performed interviews with platform experts and conducted a case study. Our intention was to validate both the pattern language and the decision model and to investigate the decision making process in this context. The results of these research steps were a refined pattern language and a reusable architectural decision model. In addition, we propose a model capturing the four levels of decision making for platform integration, as identified in our case study. Based on this model, we documented design decisions for all levels and stages of decision making.

The evidence presented in this paper indicates that the notion of decision levels and decision stages applies to architectural decision making in general. As for platform-driven software development, the four decision levels (e.g.,

architecture, platform, integration, application; see Section IV-D) could be characteristic and are likely to apply to other, platform-like software development approaches (e.g., software product lines). This motivates us to follow up on this conjecture in our future research. Our finding of multiple levels and multiple stages in architecture decision making also relates to the way architectural decision-making techniques are designed and how these decision-making techniques are classified [34]. For example, our finding converges with the recently proposed decision-making technique by Zimmermann et al. [32], also indicating the need for integrating levelled decision making. In systematic reviews on decision-making techniques (such as [34]), levelled and staged decision making have not yet been considered, for example, in terms of discriminating between domain-specific groups of decision makers (as represented by our four levels, e.g., integration architect, platform supplier).

As for our reusable, pattern-based architectural decision model (ADM), on the one hand, we plan to refine it further by performing further qualitative studies (e.g., interviews, case studies) with an increased number of participants. On the other hand, and based on the refined ADM, we want to assess its cost-benefit balance [35] for documenting architectural design rationale empirically. For example, while benefits from pattern-based decision documentation have been indicated [6–8], the use of patterns can also cause extra pattern learning and pattern search efforts.

Acknowledgements: This work was partially supported by the EU FP7 project INDENICA, grant no. 257483.

REFERENCES

- [1] R. N. Taylor and A. van der Hoek, "Software Design and Architecture: The once and future focus of software engineering," in *Proc. 2007 Future of Softw. Eng. Conf.* IEEE, 2007, pp. 226–243.
- [2] A. Jansen and J. Bosch, "Software Architecture as a Set of Architectural Design Decisions," in *Proc. 5th Work. IEEE/IFIP Conf. Softw. Architecture.* IEEE, 2005, pp. 109–120.
- [3] J. Tyree and A. Ackerman, "Architecture Decisions: Demystifying Architecture," *IEEE Softw.*, vol. 22, no. 19-27, 2005.
- [4] O. Zimmermann, T. Gschwind, J. Kuester, F. Leymann, and N. Schuster, "Reusable Architectural Decision Models for Enterprise Application Development," in *Proc. 3rd Int. Conf. Quality of Softw. Architecture*, ser. LNCS, vol. 4880. Springer, Jul. 2007, pp. 15–32.
- [5] Y. Ghanam, F. Maurer, and P. Abrahamsson, "Making the leap to a software platform strategy: Issues and challenges," *Inform. Software Tech.*, vol. 54, no. 9, pp. 968–984, 2012.
- [6] N. Harrison, P. Avgeriou, and U. Zdun, "Using Patterns to Capture Architectural Decisions," *IEEE Softw.*, vol. 24, no. 4, pp. 38–45, Jul. 2007.
- [7] O. Zimmermann, U. Zdun, T. Gschwind, and F. Leymann, "Combining Pattern Languages and Reusable Architectural Decision Models into a Comprehensive and Comprehensible Design Method," in *Proc. 7th Work. IEEE/IFIP Conf. Softw. Architecture.* IEEE, 2008, pp. 157–166.
- [8] U. van Heesch and P. Avgeriou, "A Pattern-based Approach Against Architectural Knowledge Vaporization," in *Proc. 14th Annu. Europ. Conf. on Pattern Languages of Programs*, ser. CEUR Workshop Proceedings, vol. 566. CEUR-WS.org, 2009.
- [9] B. Kitchenham, O. Brereton, D. Budgen, M. Turner, J. Bailey, and S. Linkman, "Systematic literature reviews in software engineering – a systematic literature review," *Inform. Software Tech.*, vol. 51, no. 1, pp. 7–15, 2009.
- [10] C. B. Seaman, "Qualitative methods in empirical studies of software engineering," *IEEE Trans. Softw. Eng.*, vol. 25, no. 4, pp. 557–572, Jul. 1999.
- [11] C. Robson, *Real World Research - A Resource for Social Scientists and Practitioner-Researchers.* Blackwell Publishing, 2002.
- [12] I. Lytra, S. Sobernig, H. Tran, and U. Zdun, "A Pattern Language for Service-Based Platform Integration and Adaptation," in *Proc. 17th Annu. Europ. Conf. on Pattern Languages of Programs.* Hillside, Jul. 2012.
- [13] A. Tashakkori and C. Teddlie, *Handbook of Mixed Methods in Social & Behavioral Research*, 2nd ed. Sage Publications, Inc., 2010.
- [14] F. Buschmann, K. Henney, and D. C. Schmidt, *Pattern-Oriented Software Architecture – A Pattern Language for Distributed Computing.* Wiley, 2007.
- [15] M. Fowler, *Patterns of Enterprise Application Architecture.* Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc., 2003.
- [16] G. Hohpe and B. Woolf, *Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions.* Addison-Wesley, 2004.
- [17] M. Völter, M. Kircher, and U. Zdun, *Remoting Patterns: Foundations of Enterprise, Internet and Realtime Distributed Object Middleware.* Wiley, 2005.
- [18] R. Daigneau, *Service Design Patterns: Fundamental Design Solutions for SOAP/WSDL and Restful Web Services.* Addison-Wesley, 2012.
- [19] C. Hentrich and U. Zdun, *Process-Driven SOA: Patterns for Aligning Business and IT.* Infosys Press, 2012.
- [20] E. Gamma, R. Helm, R. E. Johnson, and J. Vlissides, *Design Patterns – Elements of Reusable Object-Oriented Software.* Addison-Wesley, 1994.
- [21] F. Buschmann, R. Meunier, H. Rohnert, P. Sommerlad, and M. Stal, Eds., *Pattern-Oriented Software Architecture – A System of Patterns.* Wiley, 2000.
- [22] P. Avgeriou and U. Zdun, "Architectural patterns revisited – A pattern language," in *Proceedings of EuroPLoP 2005*, Irsee, Germany, Jul. 2005, pp. 1–39.
- [23] F. Buschmann, K. Henney, and D. C. Schmidt, *Pattern-Oriented Software Architecture – On Patterns and Pattern Languages.* Wiley, Apr. 2007.
- [24] J. O. Coplien, *Software Patterns.* SIGS Books, 1996.
- [25] I. Lytra, S. Sobernig, and U. Zdun, "Architectural Decision Making for Service-Based Platform Integration: A Qualitative Multi-Method Study," Faculty of Computer Science, University of Vienna, Tech. Rep. TR-SWA-20120601, 2012.
- [26] S. E. Hove and B. Anda, "Experiences from Conducting Semi-structured Interviews in Empirical Software Engineering Research," in *Proc. 11th IEEE Int. Software Metrics Symp.* IEEE, 2005, pp. 23–32.
- [27] B. Glaser and A. Strauss, *The Discovery of Grounded Theory.* Aldin, 1967.
- [28] K. Czarniecki, S. Helsen, and U. W. Eisenecker, "Staged Configuration Through Specialization and Multi-Level – Configuration of Feature Models," *Softw. Process: Improvement & Practice*, vol. 10, no. 2, pp. 143–169, 2005.
- [29] C. Wohlin, P. Runeson, M. Host, C. Ohlsson, B. Regnell, and A. Wesslén, *Experimentation in Software Engineering – An Introduction.* Kluwer Academic Publishers, 2000.
- [30] P. Brereton, B. Kitchenham, D. Budgen, M. Turner, and M. Khalil, "Lessons from applying the systematic literature review process within the software engineering domain," *J. of Syst. and Softw.*, vol. 80, no. 4, pp. 571–583, 2007.
- [31] D. Falessi, M. Babar, G. Cantone, and P. Kruchten, "Applying empirical software engineering to software architecture: challenges and lessons learned," *Emp. Softw. Eng.*, vol. 15, pp. 250–276, 2010.
- [32] O. Zimmermann, J. Koehler, F. Leymann, R. Polley, and N. Schuster, "Managing architectural decision models with dependency relations, integrity constraints, and production rules," *J. of Syst. and Softw.*, vol. 82, no. 8, pp. 1249–1267, 2009.
- [33] R. Capilla, O. Zimmermann, U. Zdun, P. Avgeriou, and J. Küster, "An Enhanced Architectural Knowledge Metamodel Linking Architectural Design Decisions to other Artifacts in the Software Engineering Lifecycle," in *Proc. 5th Europ. Conf. Softw. Architecture*, ser. LNCS, vol. 6903. Springer, 2011, pp. 303–318.
- [34] D. Falessi, G. Cantone, R. Kazman, and P. Kruchten, "Decision-making techniques for software architecture design: A comparative survey," *ACM Comput. Surv.*, vol. 43, no. 4, pp. 33:1–33:28, 2011.
- [35] D. Falessi, L. Briand, G. Cantone, R. Capilla, and P. Kruchten, "The Value of Design Rationale Information," *ACM Trans. Softw. Eng. Method.*, 2012, to be published.