

Managing Architectural Decision Models with Dependency Relations, Integrity Constraints, and Production Rules

Olaf Zimmermann¹, Jana Koehler¹,
Frank Leymann², Ronny Polley¹, Nelly Schuster¹

¹ IBM Research GmbH
Zurich Research Laboratory, Säumerstrasse 4, 8803 Rüschlikon, Switzerland
{olz, koe, rpo, nes}@zurich.ibm.com

² Universität Stuttgart, Institute of Architecture of Application Systems
Universitätsstraße 38, 70569 Stuttgart, Germany
frank.leymann@iaas.uni-stuttgart.de

Abstract. Software architects consider capturing and sharing architectural decision knowledge increasingly important; many tacit dependencies exist in this knowledge. Architectural decision modeling makes these dependencies explicit and serves as a foundation for decision knowledge management tools. In practice, however, text templates and informal rich pictures rather than models are used to capture the decision knowledge; a formal definition of model entities and their relations is missing in the current state of the art. In this paper, we propose such a formal definition of architectural decision models as directed acyclic graphs with several types of nodes and edges. In our models, architectural decision topic groups, issues, alternatives, and outcomes form trees of nodes connected by edges expressing containment and refinement, decomposition, and triggers dependencies, as well as logical relations such as (in)compatibility of alternatives. The formalization can be used to verify integrity constraints and to organize the decision making process; decision types and dependency patterns can be identified. A reusable architectural decision model supporting service-oriented architecture design serves as our example. We present tool support and give a quantitative evaluation.

Keywords: Architectural decision, decision type, dependency pattern, dependency relation, graph, integrity constraint, knowledge management, model, design decision rationale, SOA, Web services, UML

1 Introduction

Having been neglected both in academia and industry for a long time, the importance of *architectural decision capturing and sharing* is now widely acknowledged [4][12][21]. However, existing work focuses on capturing and visualizing decisions that have been made already. In practice, text templates and informal rich pictures are used to capture this knowledge. Little emphasis is spent on specifying the

dependencies between decisions and on sharing information about architectural decisions required and alternatives available. Lack of decision capturing rigor is a possible source of quality problems with the software architectures under construction; insufficient incentives, methods, and tools for decision sharing inhibit a reuse of knowledge and an exchange between practitioners on different projects.

Existing decision capturing and sharing models and tools lack a formalization of decisions and their dependencies. Extending our existing Unified Modeling Language (UML) domain metamodel [25], we apply set and graph theory concepts in this paper to formally define architectural decision issues, alternatives, and outcomes and several types of containment and dependency relations such as decomposition, refinement, triggers, forces, and (in)compatibility. We use these logical and temporal relations to structure the decision models and to order the decision making process.

Such formalization of the data structures in an architectural decision model is useful for several other purposes. It allows knowledge engineers to *verify the quality* of a reusable decision model developed in a practitioner community. Software architects can *evaluate* the models created on individual projects. A decision order makes the models comparable. Graph traversal algorithms can be developed, e.g., calculating path lengths in support of *model maintenance*. Decision types and dependency patterns can also be defined, which helps to detect the incompleteness or inconsistency of a decision model. Finally, knowledge engineers working in other decision capturing domains, e.g., not SOA, or not even software architecture, can *reuse the model structure* to organize their knowledge.

The remainder of this paper is structured in the following way. Section 2 discusses related work and examples from Service-Oriented Architecture (SOA) design. Section 3 introduces our UML domain metamodel and types of decisions we observed to recur in enterprise application development and SOA design. Section 4 presents the formalization of architectural decision models. Section 5 introduces decision dependency patterns. Section 6 discusses how we validated model, content, and implementation. Section 7 concludes with a summary and an outlook to future work.

2 Related Work

Bass et al. mention the term *architectural decision*, but not fully define it in “Software Architecture in Practice” from 1998 [3]. Kruchten et al. [12] define an ontology that describes the attributes that should be captured for a decision, the types of decisions (e.g., executive, existence, and property decisions), when and how decisions are made (i.e., their lifecycle), and several types of decision dependencies. They also focus on the visualization of the decisions and identify many use cases for decision knowledge. Their ontology is semantically rich and defines the knowledge domain both broadly and deeply. However, it is described informally only (i.e., in text). Moreover, design problem and solution are treated as one entity (i.e., alternatives are a dependency type, not an entity). Hence, decisions required (which we refer to as issues) and decisions made (which we refer to as outcomes) can not be separated, which limits the reusability of the modeled knowledge. Finally, there are no concepts for structuring and ordering decision models apart from the decision types.

Jansen and Bosch [5] view software architecture as a composition of a set of design decisions. They make the case for decisions to be a first class architecture design concept. Their model for architectural design decisions focuses on change over time as a dominating force driving the decision making. In their metamodel, they distinguish design problems and solutions to them, and outline the attributes that are required to capture related knowledge. Design fragments make it possible to integrate decision models with models for other viewpoints (e.g., logical components and connectors). The metamodel is introduced in text and figures; dependencies between different problems or different solutions remain implicit (i.e., decisions depend on each other if they deal with the same or with related design fragments). There is no overarching model structuring scheme. The reuse of architectural decision knowledge is only touched upon: Design patterns are mentioned as a source of reusable solutions.

Several other decision capturing templates exist in industry and academia, which can also be viewed as informally specified metamodels. For instance, the IBM Unified Method Framework (UMF) defines such template in its “architectural decisions” artifact. Architects’ Workbench (AWB) provides modeling tool support for this and many other UMF artifacts [1]. The “Group ADs by Topic” viewpoint in AWB introduces a topic hierarchy and defines an outcome attribute in the decision entity; alternatives are modeled as a separate entity. UMF was formerly known as IBM Global Services Method and has been in use on professional services engagements for IBM clients since 1998. One of the IBM reference architectures comes with a filled out architectural decisions artifact, which contains architectural decisions required during Web application design. Having worked with this artifact, Tyree and Akerman [21] defined another rich decision capturing template, structured into 13 sections. Later on, they proposed an entire ontology to support the design of software architectures [2]. SEURAT, PAKME, ADDSS, The Knowledge Architect, AREL, Archium, and ArchStudio are additional tools providing decision modeling capabilities and supporting metamodels. An overview of these tools is given in [5].

In the patterns community, several schools of thought and many pattern templates exist, which can also be used to capture architectural decisions [8]. Requirements in areas such as performance and extensibility typically are referenced in textual *intent* or *forces* sections. Many pattern languages remain on an abstract, conceptual level; others specialize on a single problem or technology domain such as *enterprise application architecture* [6] or *process-driven SOA* [23]. Patterns for process-driven SOA describe how to automate the management of long-running business processes such as loan approval processing or order management along supply chains (problem domain) with workflow engines and integration middleware (technology domain). The activity flow in such processes can be specified using Business Process Modeling (BPM) tools and implemented as a network of communicating *Web service consumers and providers*. In our earlier work, we demonstrated that the relationship between patterns languages and architectural decision models is synergetic: We position architectural patterns as conceptual architecture alternatives in our *reusable* architectural decision models [25][27]. The pattern texts serve as source of architectural decision knowledge

Defining templates or metamodels and referencing patterns is a good starting point towards more systematic and rigorous decision capturing; however, it does not remove real-world inhibitors for sustainable and maintainable decision sharing such as no immediate benefits, budget and scheduling problems, and lack of tools. Tang et

al. report these and several more inhibitors in [20]. To address these issues, we formalize the concepts in existing metamodels and templates and extend them with support for reuse and collaboration: If a comprehensive architectural decision model is created for a certain domain, which can be tailored for particular project at project initiation time, the benefits reported in the literature can be realized by a community of architects over a longer period of time. The budget and tools issues are then faced by an organizational unit (e.g., architecture management group in an enterprise or a community of practice in a professional services firm) rather than by individuals or by project teams. Hence, there are better chances for overcoming them. For instance, a knowledge engineer can be tasked with the creation of a reusable architectural decision model, which is then used by architects on multiple projects.

In the next section, we introduce such reusable architectural decision model and an underlying metamodel. In Section 6, this decision model is presented in more detail.

3 A Domain Metamodel for Capturing Architectural Decisions

A domain metamodel for architectural decision capturing must be expressive enough to support the use cases from [12]. In our reuse and collaboration context, additional use cases are education, knowledge exchange, design method support, review technique, and governance instrument. The metamodel should only define a small set of mandatory attributes so that practitioners are not overwhelmed with information when populating and studying decision models. The metamodel must support reuse and collaboration. It must be machine readable and translatable into other specifications, e.g., into Web services contracts and relational database schemas, so that tool support for decision modeling and dependency management can be built.

In our earlier work, we derived such a metamodel from earlier proposals [1][4][5][21] and our practical industry experience [1], and defined three processing steps, decision identification, decision making, and decision enforcement [25]. An updated version of that model is shown in Figure 1. It uses Unified Modeling Language (UML) classes to introduce the three core entities *ADIssue*, *ADAlternative*, and *ADOutcome*; *ADTopicGroups* and *ADLevels* are supplemental structuring concepts.¹

An *ADIssue* instance informs the architect that a single architecture design problem has to be solved. *ADAlternative* instances then present possible solutions to this problem. Finally, *ADOutcome* instances record an actual decision made to solve the problem including its rationale. *ADTopicGroup* instances bundle related issues.

We distinguish decisions made and decisions required to facilitate reuse: *ADIssue* and *ADAlternative* provide reusable background information about decisions required to the architect: The *problemStatement* characterizes an *ADIssue* on an introductory level, while *backgroundReading* and *knownUses* (*ADAlternative*) point to further information. The *decisionDrivers* attribute states types of non-functional requirements, including software quality attributes and environmental issues such as budget and skill availability; the patterns community uses the term *forces* synonymously. The *responsible* (role) and *phase* attribute serve as link to general-purpose methodologies

¹ In our previous work, the *ADIssue* was called *AD*; *ADTopicGroup* was called *ADTopic*.

such as the Rational Unified Process (RUP) [11]. ADIssues and ADAAlternatives capture project-independent architectural decision knowledge.

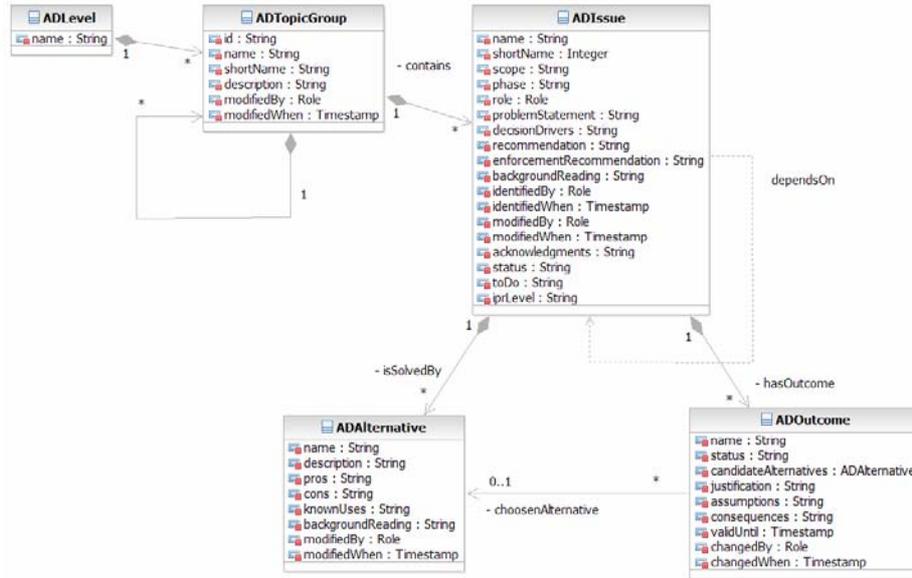


Fig. 1. UML metamodel for architectural decision capturing and reuse.

ADOutcome instances capture project-specific knowledge about decisions made. The *justification* information refers to actual requirements (“sub-second response time in customer interface”), as opposed to the ADIssue-level decision drivers which only list types of requirements (“performance, i.e., response time and throughput”). These two aspects of the knowledge have different reuse characteristics: the background information has even more reuse potential than the project-specific rationale. A second reason for factoring out ADOOutcome as a separate entity is that the same ADIssue might pertain to many elements in a design model, e.g., business processes and Web service operations in SOA. Therefore, *types* of design model elements are referenced via the *scope* attribute in the ADIssue. ADOOutcome instances then can be created dynamically on projects, and can refer to design model element *instances* via their *name*.

To give an example, a business process model for order management might state that three “customer enquiry”, “claim check”, and “risk assessment” business processes have to be implemented in an insurance industry case.² One ADIssue is to select an INTEGRATION TECHNOLOGY to let the activities in each of the three business processes interact with other systems, with ADAAlternatives such as WEB SERVICES and RESTFUL INTEGRATION.³ Problem statement (“Which technology should be used to let the activities in the business process communicate with Web services and legacy systems?”) and decision drivers (“interoperability”, “reliability”, and “tool support”) are the same for all three business processes. Hence, it is sufficient to create a single

² See [1] for an order management SOA case study from the telecommunications industry.

³ In all further examples, we set ADIssues and ADAAlternatives in THIS FONT (small caps).

ADIssue instance which has a “business process” scope. This value of the scope attribute refers to a type of SOA-specific design model element.

Decision outcome information such as the chosen alternative and its justification depends on the individual requirements of each process, e.g., “for customer enquiry, we decide for WEB SERVICES as Java and C# components have to be integrated in an interoperable and reliable manner” and “for risk assessment, we select RESTFUL INTEGRATION because not all of the involved backend systems provide a SOAP message interface described by a WSDL contract”. Hence, three ADOutcome instances are created and associated with the same ADIssue. These instances capture the process-specific decision and its rationale. They refer to the actual business processes in their name attributes (“customer enquiry”, “claim check”, and “risk assessment”).

Closely related ADIssues are grouped into *ADTopicGroups*, which form a hierarchy. This hierarchy serves as a table of content: Each ADTopicGroup hierarchy is assigned to one of several *ADLevels* of refinement, e.g., *conceptual level*, *technology level*, or (*vendor*) *asset level*. This structure is motivated by our observation that when designing enterprise applications, the technical discussions often circle around detailed features of certain vendor products, or the pros and cons of specific technologies, whereas many highly important strategic decisions and general concerns are underemphasized. These discussions are related, but should not be merged into one; they reside on different refinement levels. Separating design concerns in such a way is good practice; Fowler [6] and RUP with its elaboration points recommend a similar incremental approach for UML class diagrams used as design models. We adopted this recommendation for decision models and made the three refinement levels explicit in our domain metamodel. It is possible to select other ADTopicGroup hierarchies. For instance, panes in enterprise architecture frameworks and logical viewpoints can also be used as structuring mechanisms.

Decision dependencies are explicitly modeled as UML associations between ADIssues. We defined a single *dependsOn* dependency type in Figure 1; in Section 4, we introduce additional dependency types that correspond to those defined in [12].

Rationale. Our metamodel extends that from [1] and [5], e.g., with the levels concept. Jansen and Bosch also separate problem (issue) from solution (alternative), and define how to scope decisions via design fragments. Similar entities and concepts for method alignment can be found in the core model defined by de Boer et al. [4], which was developed independently of and simultaneously to our UML model. Unlike de Boer et al., we also define attributes, which is required to support reuse and collaboration as motivated in Section 1 and Section 2. In particular, we define attributes that are required for lifecycle management of ADIssues in the reusable part (e.g., *role*) and ADOutcomes in project-specific decision models (e.g., *changedBy*).

Example. A *Reusable Architectural Decision Model (RADM) for SOA* serves as a running example throughout this paper. It was created in an industrial decision mining project [26] that started in 2006 (see Section 6 for more information). All 389 decisions captured so far conform to the metamodel shown in Figure 1.

Table 1 shows several ADIssue examples from the RADM for SOA and assigns them to seven decision types. We found many instances of these seven decision types during the creation of the RADM for SOA (see Section 6 and [25] for rationale):

Table 1. Decision types and SOA examples (RADM for SOA)

Decision type	ADLevel	ADTopicGroups and ADISSUES (in SMALL CAPS)
<i>Executive decisions, requirements analysis decisions</i>	Executive level	Out of scope of this paper, introduced in [25] and elaborated upon in [27].
<i>Pattern Selection Decisions (PSDs)</i>	Conceptual level	ADTopicGroup Service Layer Realization Decisions: IN MESSAGE GRANULARITY MESSAGE EXCHANGE PATTERN INVOCATION TRANSACTIONALITY PATTERN ADTopicGroup Process Layer Realization Decisions: SERVICE COMPOSITION PARADIGM MACROFLOW MICROFLOW ADTopicGroup Integration Layer Realization Decisions INTEGRATION STYLE BROKER, ADAPTER, REGISTRY
<i>Pattern Adoption Decisions (PADs)</i>	Conceptual level	Process Layer Realization Decisions: MACROFLOW MICROFLOW PROCESS ACTIVITY TRANSACTIONALITY (PAT) Integration Layer Realization Decisions: COMMUNICATIONS TRANSACTIONALITY (CT) REGISTRY LOOKUP TIME
<i>Technology Selection Decisions (TSDs)</i>	Technology level	ADTopicGroup Service Layer Technologies: TRANSPORT PROTOCOL CHOICE MESSAGE EXCHANGE FORMAT ADTopicGroup Process Layer Technologies: SERVICE COMPOSITION LANGUAGE ADTopicGroup Integration Layer Technologies: INTEGRATION TECHNOLOGY AUTHORIZATION TECHNOLOGY
<i>Technology Profiling Decisions (TPDs)</i>	Technology level	Service Layer Technologies: SOAP COMMUNICATION STYLE WEB SERVICES TRANSACTIONALITY Process Layer Technologies: BPEL VERSION COMPENSATION TECHNOLOGY Integration Layer Technologies: TRANSPORT QOS XML SCHEMA CONSTRUCTS
<i>Vendor Asset Selection Decisions (ASDs)</i>	Vendor asset level	ADTopicGroup Service Layer Assets: SOAP ENGINE ADTopicGroup Process Layer Assets: BPEL ENGINE SCA IMPLEMENTATION ADTopicGroup Integration Layer Assets: ESB GATEWAY
<i>Vendor Asset Configuration Decisions (ACDs)</i>	Vendor asset level	Service Layer Assets: IBM/AXIS SOAP ENGINE DEPLOYMENT MODE AXIS SOAP ENGINE DEPLOYMENT MODE Process Layer Assets: (WPS BPEL) INVOKE ACTIVITY TRANSACTIONALITY SCA QUALIFIERS Integration Layer Assets: ESB TOPOLOGY

In addition to the three levels already introduced, we use an *executive level* in the RADM for SOA, which comprises executive decisions as defined in the taxonomy from Kruchten et al.

Pattern Selection Decisions (PSDs) are concerned with choosing certain architectural patterns from the vast body of patterns available in the literature. Pattern Adoption Decisions (PADs) also deal with architecture and design patterns, but in a more detailed way, e.g., selecting certain pattern variants and pattern primitives once a PSD has been made. Such PADs often can be found in the pattern texts, e.g., in bulleted lists, cheat sheets and overview diagrams in patterns books. Pattern language primitives and grammars as defined by Zdun et al. [23] are another source of PAD identification.

Technology Selection Decisions (TSDs) select certain technologies that implement the selected and adopted patterns; Technology Profiling Decisions (TPDs) follow them, specifying implementation details such as subsets of technology standards to be employed. An example TPD is the decision about the XML SCHEMA CONSTRUCTS that are selected from the many options in the CML schema standard to serve as request and response message parameters defined for service operations.

Asset Selection Decisions (ASDs) pick commercial products or open source assets supporting the selected and profiled technologies; Asset Configuration Decisions (ACDs) then cover implementation and deployment details of these products.

Let us give another, more advanced example. When implementing the three business processes for order management introduced above, a conceptual PSD for a SERVICE COMPOSITION PARADIGM is required, deciding whether the processes should be made executable in a workflow engine, or be realized in traditional programming language code. If a workflow engine is decided for, a related TSD is to agree on a SERVICE COMPOSITION LANGUAGE such as Business Process Execution Language (BPEL). Another related issue is to select a BPEL ENGINE as an ASD, e.g., Active BPEL, IBM WebSphere Process Server or Oracle BPEL Process Manager. For each of the activities in a business process and for each invoked Web service, the INVOCATION TRANSACTIONALITY PATTERN and INTEGRATION STYLE have to be decided. These issues have several related PADs, TPDs, and ACDs. This fairly complex set of issues will serve as our example in this paper.

While the content in this particular RADM is specific to enterprise application development and SOA, the concepts presented in the next section provide generic solutions to the decision capturing and sharing problems outlined in Sections 1 and 2.

4 A Formal Model for Decision Modeling with Reuse

Existing decision capturing approaches are based on text templates or informally specified metamodels. Their main usage scenario, architecture documentation, has a retrospective nature. In such a setting, each decision is captured from scratch and ad hoc as it is made during design. In some approaches, it is mined from other artifacts. On the contrary, our approach emphasizes the proactive sharing of reusable background information about recurring design issues, captured in ADIssue and ADAAlternative instances. Such reusable decision model can steer software architects

through the decision making, informing them about decisions required and documenting the problems to be solved. For recurring ADIssues, only the ADOutcome instances have to be created on the project.

To be able to use a reusable decision model in such active guiding role, additional concepts are required. For instance, a structure must be defined that organizes large models and makes them consumable; a decision making order must be specified. To do so, we complement the UML model from Section 3 with formal definitions now. The rationale behind and motivating examples for each concept come from the SOA domain; however, it is an explicit design goal for our modeling concepts that the concepts can also be applied to other architectural domains.

4.1 Elementary Definitions for Architectural Decision Modeling

Basic concepts from set and graph theory are adequate to define the entities in the UML model and the relations between them. We begin with representations for the UML model elements ADTopicGroup, ADIssue, and ADAAlternative from Figure 1.

Definition 1 (Architectural Decision Topic Groups T) *Let T be a set of architectural decision topic groups $T = \{(n, s, d) \mid n, s, d \in \text{Strings}\}$ where the tuple (n, s, d) represents the name, short name, and description of an architectural decision topic group T .⁴*

Rationale and example: An architectural decision topic group (short: topic group) represents closely related design concerns. For instance, in the RADM for SOA, one topic group per architectural layer is defined on each refinement level (Table 1 in Section 3). An example is the ADTopicGroup “service layer realization decisions”. It is worth noting that our topic groups are different from the topics in [4]. They do not represent individual design issues, but group such issues. Representing individual design issues is the purpose of the next entity:

Definition 2 (Architectural Decision Issues I) *Let I be a set of architectural decision issues $I = \{(n,s,p,r, \{tt\}) \mid n, s, p, r, tt \in \text{Strings}\}$ where n is a name, s a scope, p a project phase, r a role attribute, and tt a set of topic tag strings.*

Rationale and example: An architectural decision issue (short: issue) represents a single design concern. Name, scope, phase, role are describing texts. The name is used to identify and list issues. The topic tags index the model content. This information can be used to locate issues by subject area keyword. In the RADM for SOA outlined in Section 3, we use the names of the decision types from Table 1 as topic tags, as well as important non-functional concerns such as security and transaction management. The architect can query the model for all PSDs (see Section 3), all issues dealing with security and/or transaction management, etc.

In our SOA decision model, two PSDs deal with the MESSAGE EXCHANGE PATTERN (dealing with the abstract protocol syntax and synchrony of service invocations) [9] and the INVOCATION TRANSACTIONALITY PATTERN (dealing with

⁴ The other attributes from the UML model are irrelevant for the formalization.

system transactions as an approach to protecting shared resources from invalid concurrent access, e.g., lost updates and phantom reads [6]). Another issue is the IN MESSAGE GRANULARITY PSD, which concerns the syntax (breadth and depth) of the in message parameters. These decisions are listed in Table 1 in Section 3.

An architectural decision issue captures a single design concern or problem without modeling possible solutions to it. Architectural decision alternatives do so:

Definition 3 (Architectural Decision Alternatives A) Let A be a set of architectural decision alternatives $A = \{(n, d) \mid n, d \in \text{Strings}\}$ where n is a name and d is a solution description.

Rationale and example: An architectural decision alternative (short: alternative) presents a single solution to the design issue expressed by an ADIssue. For instance, the MESSAGE EXCHANGE PATTERN can decide between synchronous REQUEST REPLY and asynchronous ONE WAY message exchange. Two alternatives for the INVOCATION TRANSACTIONALITY PATTERN might be TRANSACTION ISLANDS (do not let service consumer and provider share a single transaction context) and TRANSACTION BRIDGE (propagate transaction context with a service invocation) [24].

Definition 4 (contains relations $\prec_T, \prec_I, \prec_A, \prec$) Let $\prec_T \subseteq T \times T$ be a contains relation defined between topic groups, $\prec_I \subseteq T \times I$ be a contains relation defined between topic groups and issues, and $\prec_A \subseteq I \times A$ be a contains relation defined between issues and alternatives. Subsequently, we will only speak of the contains relation $\prec = \prec_T \cup \prec_I \cup \prec_A$. If $(a \prec b)$, we say that a contains b and b is contained in b .

Rationale and example: The contains relation \prec allows us to define a single hierarchical structure, which serves as a table of content, allowing the user to locate issues and alternatives easily in the reusable architectural knowledge and helping the knowledge engineer to avoid undesired redundancies. One or more alternatives solve a particular issue. Related issues can be put into one topic group. Related topic groups can be placed in the same parent topic group. Figure 2 illustrates the tree structure resulting from the \prec relation:

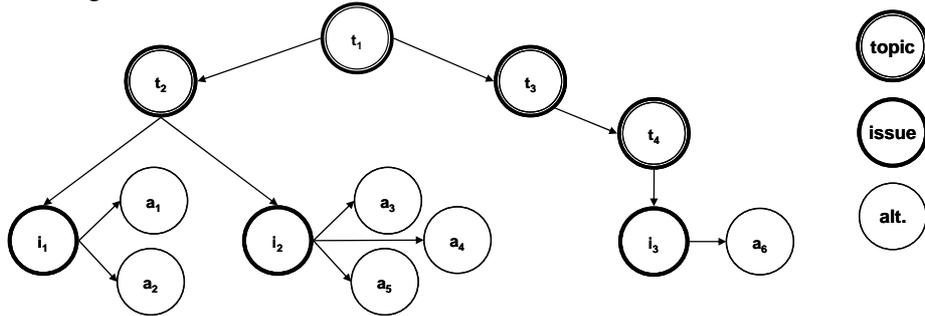


Fig. 2. General organization of an architectural decision tree, indices reflect an ordering of the topic groups, issues and alternatives.

In the UML metamodel in Figure 1, the \prec relation is represented by the three associations (arrows with filled with solid diamonds at originating end) that express

physical containment between ADTopicGroups, ADIssues and ADAAlternatives, respectively.

Unlike other proposals, we define only a single tree structure (i.e., no overlays), and one alternative can only be a solution to one design issue. This modeling decision is justified by our emphasis on reuse: In reusable architectural decision models, knowledge engineers describe the attributes of the alternatives relative to the problem statement and the decision drivers of an issue, which makes it necessary to define a 1:n relation (also see UML model in Section 3). If a pattern, technology, or asset solves multiple problems, it is referenced in multiple alternatives. We faced a tradeoff between normalization (i.e., no redundancy) and precision (accuracy) when making this modeling decision; the latter requirement is more important in our usage context.

Definition 5 (Architectural Decision Tree \mathbb{T}) Using T, I, A , and the \prec relation, we can define an architectural decision tree $\mathbb{T} = (T \cup I \cup A, \prec)$ with a single root node $t_0 \in T$ called the root topic.⁵ In \mathbb{T} , a topic group contains zero or more other topic groups and issues, while an issue may contain one or more alternatives. In this tree, each topic group $t \in T$ except the root topic is contained in exactly one other topic group $t_i \in T$:

$$\forall t, t_i, t_j \in T: (t_i \prec t) \wedge (t_j \prec t) \Rightarrow t_i = t_j$$

Each issue $i \in I$ must be contained in exactly one topic group $t \in T$:

$$\begin{aligned} & \forall i \in I \exists t \in T: (t \prec i) \\ & \forall i \in I, t_i, t_j \in T: (t_i \prec i) \wedge (t_j \prec i) \Rightarrow t_i = t_j \end{aligned}$$

Each alternative $a \in A$ must be contained in exactly one issue $i \in I$:

$$\begin{aligned} & \forall a \in A: \exists i \in I (i \prec a) \\ & \forall i_i, i_j \in I, a \in A: (i_i \prec a) \wedge (i_j \prec a) \Rightarrow i_i = i_j \end{aligned}$$

Rationale and example: Modeling architectural decisions in itself is not new: Ran and Kuusela also propose (but do not formalize) the notation of Design Decision Trees (DDTs) [16]. Our formalization allows us to define advanced concepts later.

The topic group hierarchy may mimic the containment hierarchy of a design model, e.g., beginning with architectural layers. In our RADM for SOA, the hierarchy resembles the containment hierarchy of a Web service definition. “Service” is one of the conceptual patterns that define SOA as architectural style and Web Services Description Language (WSDL) [22] is one of several technology options to express service contracts; WSDL port types define service operations through in messages accepted and out messages returned. Both service pattern and WSDL technology have several architectural decisions attached (see Table 1). Consequently, “service layer realization decisions” is a topic group on the conceptual level, which has child topic groups such as “operation design” and “message design” (not shown in Table 1). Such containment relations between design model elements exist in many application

⁵ In graph theory, a directed graph is a pair $G = (V, E)$ where V is a set of vertices (or nodes) and E is a subset of $V \times V$ relations (ordered pairs) called edges (arcs). A graph that does not contain any cycles is an acyclic graph. A directed acyclic graph is often called a DAG. A tree is a DAG with a single root node and a single path from any node to the root node.

genres and architectural styles. Architectural layering is a popular structuring principle [6].

Figure 3 instantiates the abstract tree structure from Figure 2 for parts of our SOA example:

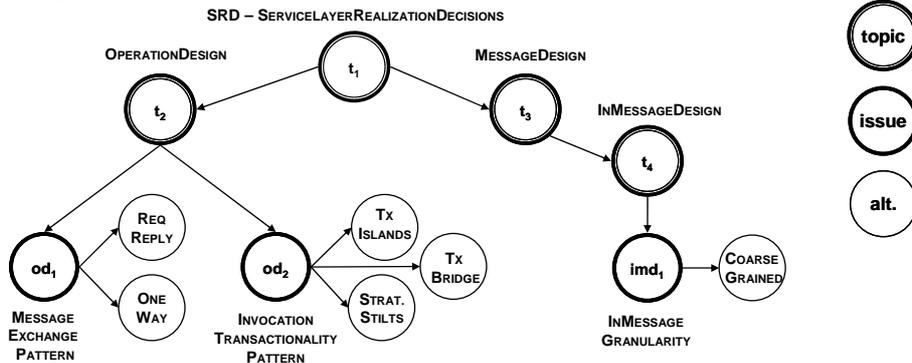


Fig. 3. An instantiated example tree showing a subset of issues that must be resolved when adding Web services to an architecture.

Definition 6 (Ordered Tree \mathbb{T}) We define an ordering among the child nodes of identical type (topic group, issue, alternative) contained in a node in order to be able to enumerate sibling nodes of the same type sharing the same parent node, i.e., we introduce $<_T, <_I, <_A$.

Rationale and example: An ordering relation defines a recommended reading sequence, and can be used to express integrity constraints on architectural decision trees (which we will define later). In the simplest case, the $<_T, <_I$, and $<_A$ relations can be the alphanumeric sorting of the topic group, issue, and alternative names. Note that a topic group may contain other topic groups and issues. In this case, we order all topic siblings before all issue siblings. This yields an ordered tree \mathbb{T} ; we refer to its total order relation as $<$.

4.2 Multi-Level Architectural Decision Model and Logical Relations

The meta model from Section 3 and the elementary definitions from Section 4.1 allow knowledge engineers to capture decisions and organize the knowledge in a topic group hierarchy. However, the resulting ordered architectural decision tree does not yet support the vision of an active, managed decision model taking a guiding role during architecture design. More relations between topics, decisions, and alternatives must be defined.⁶ In this section, we introduce logical constraints; followed by temporal dependencies in Section 4.3. Again, we apply concepts from graph theory.

Definition 7 (Architectural Decision Model \mathbb{M} , root topic, initial issue) An architectural decision model \mathbb{M} is a partially ordered set of architectural decision

⁶ Not that the UML model in Section 3 only defined a generic “dependsOn” association.

trees $T_{00}, \dots, T_{10}, \dots, T_{km}$ arranged in levels L_0, \dots, L_k . Each tree belongs to exactly one level and each level must contain at least one tree, i.e., no empty levels exist. A tree T_{ki} is the i -th tree in level k . If $k < l$, we speak of tree T_{ki} having a higher level than tree T_{lj} . Each architectural decision model \mathbb{M} has exactly one distinguished root topic, which is the root topic of T_{00} in the highest level L_0 . Accordingly, the first issue in the distinguished root topic (according to $\langle \cdot \rangle$) is identified as the initial issue.

Rationale: Architectural decision models define the multi-level structure required for knowledge bases such as that outlined in Table 1 in Section 3. The partial order assigns topics and decisions to different levels of abstraction and refinement. Figure 4 illustrates the concepts from Definition 7:

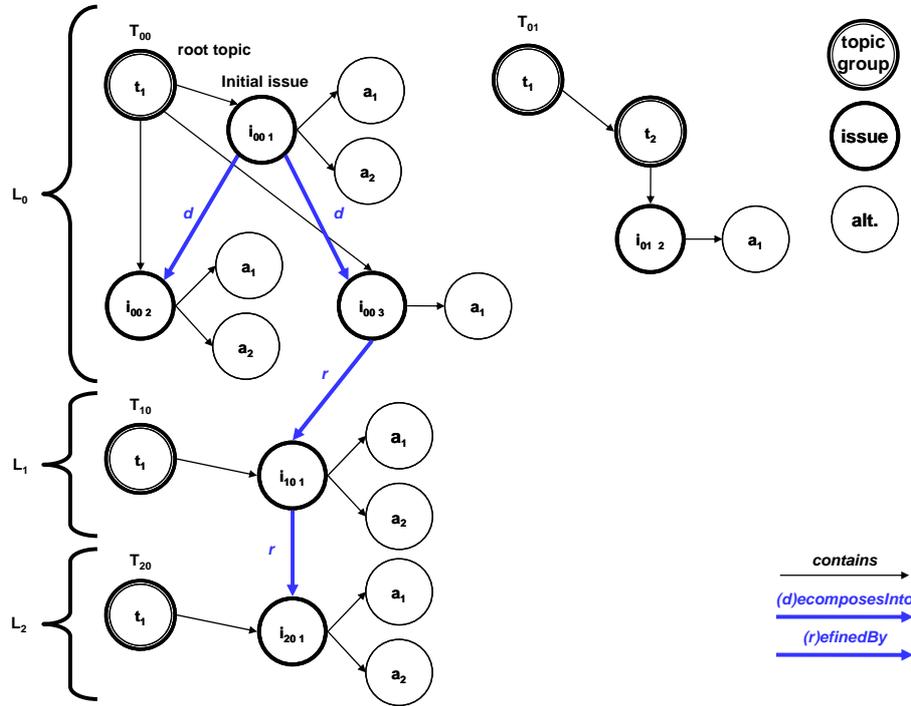


Fig. 4. A multi-level architectural decision model with four trees, root topic, and initial issue.

Figure 4 already shows relations not defined yet: i_{001} *decomposesInto* i_{002} , and i_{003} , which in turn is *refinedBy* i_{101} and then i_{201} . These relations formally capture how issues residing in different levels and trees of a model \mathbb{M} can be combined in order to express that an abstract, conceptual design is elaborated upon on the same or on a lower, more concrete level of refinement:

Definition 8 (influences, refinedBy, decomposesInto relations)

Let \mathbb{M} be an architecture decision model with levels L_0, \dots, L_k and trees T_{00}, \dots, T_{km} associated with levels $0 \dots k$. The following relations are defined between issues $i_{000}, \dots, i_{kmn} \in I$ where an issue i_{kmn} is the n -th issue in the m -th tree T_{km} contained within level L_k of a model \mathbb{M} .

- $\text{influences}(i_{j|n}, i_{k|m|o})$ with j, k, l, m, n, o arbitrary. The *influences* relation captures cross-cutting concerns between issues. It adds additional undirected edges to the model that does not necessarily have to form a connected graph. The relation is symmetric, i.e., if d_i influences d_j , then d_j also influences d_i . In addition, the *influences* relation is not reflexive, but transitive. An issue can influence several other issues and it can also be influenced by several other issues.
- $\text{refinedBy}(i_{j|n}, i_{k|m|o})$ with $j < k$ and l, m, n, o arbitrary. The *refinedBy* relation links issues that have to be investigated at several levels. It adds additional directed edges to the model that must always lead from an issue in a higher level to an issue in a lower level of the model, i.e., no cycles can occur. The relation is transitive, but not reflexive, and not symmetric. If $k = j + 1$, i.e., the refinement of an issue is contained within the subsequent level, we speak of a strict *refinedBy* relation. Issues in level 0 cannot be refinements of any other issue, while an issue in the lowest level k cannot have any refinements. If there is a relation i_1 *refinedBy* i_2 , i_1 is referred to as having an outgoing refinement relation and i_2 as having an incoming one.
- $\text{decomposesInto}(i_{j|n}, i_{k|m|o})$ with $j = k$ and l, m, n, o arbitrary. The *decomposesInto* relation expresses functional aggregation. It adds additional directed edges between issues within the same level. The relation is transitive, but neither reflexive nor symmetric. No cycles are permitted.

Table 2 summarizes the main properties of the relations.

Table 2. Decision relations between architectural decisions and their properties

Relation	Set(s)	Reflexive/ symmetric/ transitive	Cardinality (function?)	Other properties
<i>influences</i>	$I \times I$	no/yes/yes	n:m (no)	–
<i>refinedBy</i>	$I \times I$	no/no/yes	0..1:0..1 (yes)	Introduces one or more additional DAGs (i.e., no cycles permitted); only from higher to lower level (next lower if strict)
<i>decomposesInto</i>	$I \times I$	no/no/yes	0..1:n (no)	No cycles permitted. Only within same level.

Rationale and examples: We compare these relations with those defined by Kruchten et al. in Section 6. SOA examples are given later in this section (Figure 5).

The *influences* relation can be used to express cross-cutting concerns without making any assumptions about the level and order of the related decisions. For instance, the choice of a BPEL ENGINE also has to do with the AUTHORIZATION TECHNOLOGY. However, the relation is not refinement because the two issues belong to the same refinement level. It is not a decomposition either because the issues deal with different subject areas (workflow and security). The *influences* relation is often used in rapid decision capturing efforts and replaced by one of the more elaborate

forms such as *refinedBy* and *decomposesInto* as the decision model matures during subsequent knowledge engineering iterations.

The *refinedBy* relation allows us to model that the same issue typically has to be investigated at several stages of a software development process. A level can correspond to a Model-Driven Architecture (MDA) model type such as platform-independent model and platform-specific model, or to a development milestone, e.g., an elaboration point defined in RUP. A conceptual pattern such as SERVICE COMPOSITION PARADIGM abstracts away from any particular technology. Consequently, a SERVICE COMPOSITION LANGUAGE like BPEL has to be selected in refinement of the conceptual decision to adopt the WORKFLOW pattern as the SERVICE COMPOSITION PARADIGM. A particular BPEL ENGINE vendor asset has to be selected if BPEL is the selected SERVICE COMPOSITION LANGUAGE.

The *decomposesInto* relation expresses functional aggregation. When following the separation of concerns principle, complex design problems are often broken down into to smaller, more manageable units of design work (often referred to as divide-and-conquer approach to problem solving). These units can then be investigated separately (but being aware of the dependency between them).

With these relations introduced, we can define two logical constraints on architectural decision models \mathbb{M} .

Integrity Constraint 1 *The refinedBy and decomposesInto relations are mutually exclusive.*

$$\forall i_i, i_j : i_i \text{ refinedBy } i_j \Rightarrow \neg (i_i \text{ decomposesInto } i_j) \\ \text{and } \forall i_i, i_j : i_i \text{ decomposesInto } i_j \Rightarrow \neg (i_i \text{ refinedBy } i_j)$$

Rationale: This follows from our basic definitions, because the *refinedBy* relation is defined between issues residing on different levels, while the *decomposesInto* relation is only defined between issues on the same level.

Integrity Constraint 2 *If two issues have a refinedBy or a decomposesInto relation they cannot have an influences relation and vice versa.*

$$\forall i_i, i_j : i_i \text{ refinedBy } i_j \vee i_i \text{ decomposesInto } i_j \Rightarrow \neg (i_i \text{ influences } i_j) \\ \forall i_i, i_j : i_i \text{ influences } i_j \Rightarrow \neg (i_i \text{ refinedBy } i_j \vee i_i \text{ decomposesInto } i_j)$$

Rationale and example: This constraint avoids unnecessary redundancies in the model. Figure 5 adds the three levels we introduced in Section 3 to our running example, the design of transactional workflows in SOA. The topic group hierarchy is shown: three SOA layers, the service layer, the process layer, and the integration layer, are represented by separate topic groups. The PSD INVOCATION TRANSACTIONALITY PATTERN (ITP) is an example for the decomposition of a complex conceptual decision into two more primitive ones residing on the same level (here: conceptual); the related issues appear in different topic groups. The transactionality of a service operation in the SOA decision model is a non-functional design concern. It affects design model elements in the service, process, and integration layers; therefore, the PSD has relations with issues in the topic groups

representing two other SOA layers, PROCESS ACTIVITY TRANSACTIONALITY (PAT) and COMMUNICATIONS TRANSACTIONALITY (CT). PAT is an issue that resides on the process layer, CT on the integration layer. Furthermore, there are two examples of *refinedBy* relations: A strict one runs from the conceptual to the technology level (outgoing issue: CT, incoming issue: TRANSPORT QoS). Another one goes from the conceptual to the vendor asset level: The outgoing issue is PAT, the incoming is INVOKE ACTIVITY TRANSACTIONALITY (IAT).⁸

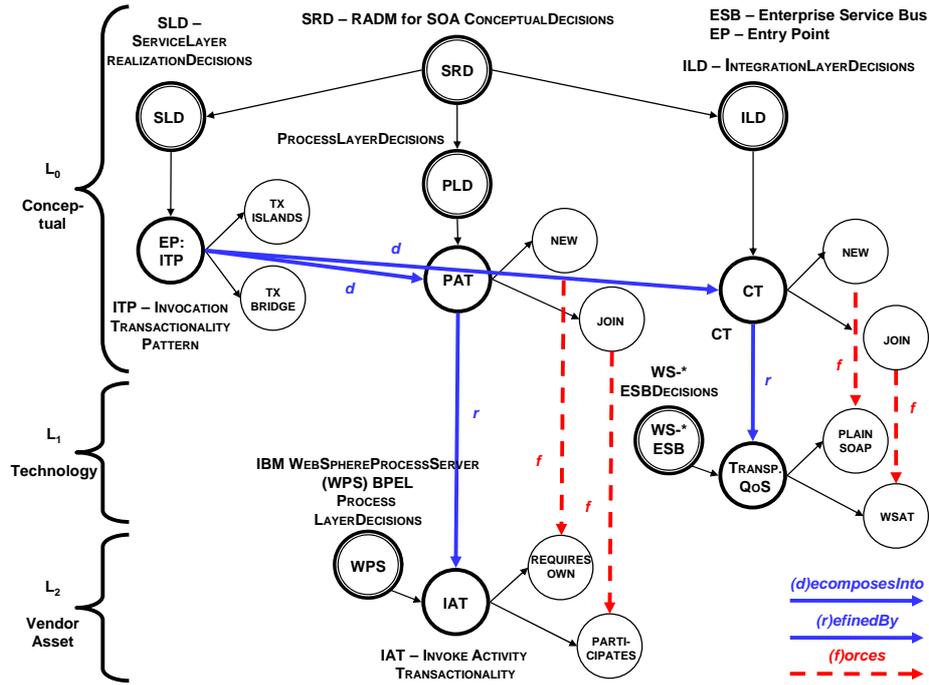


Fig. 5. Sample architectural decision model with *decomposesInto* and *refinedBy* relations.

Figure 5 also introduces a new type of relation, *forces*, expressing that certain alternatives for the conceptual decisions PAT and CT mandate the alternatives for the refining decisions on lower levels. This is one of three relations to be defined next, formally capturing the relationships that may exist between architectural decision alternatives.

Definition 9 (*forces, isIncompatibleWith, isCompatibleWith* relations) Let \mathbb{M} be an architectural decision model. Let a_i, a_k be architectural decision alternatives within \mathbb{M} . Several relations can be defined between alternatives within the same or across different levels and trees of the model.

- *forces*(a_i, a_k) with $i \neq k$ and $i_i < a_i, i_k < a_k$ implies $i_i \neq i_k$. The *forces* relation expresses that selecting an alternative a_i in one issue necessarily

⁸ This must be handled on the vendor asset level because transactionality of invoke activities is not specified by the BPEL technology standard. For details, we refer the reader to [24].

means to select an alternative a_k in another issue. It adds additional directed edges between alternatives. The relation is not reflexive and not symmetric, but transitive. It must not form any cycles.

- $isIncompatibleWith(a_i, a_k)$ with $i \neq k$. The $isIncompatibleWith$ relation expresses that certain combinations of alternatives do not work together. It adds additional undirected edges to the graph. The relation is symmetric, but neither transitive nor reflexive. It must not form any cycles.
- $isCompatibleWith(a_i, a_k)$ with i, k arbitrary. The $isCompatibleWith$ relation expresses that certain combinations of alternatives work together. The relation defines an equivalence relation, i.e., it is reflexive, symmetric, and transitive and thus identifies classes of compatible decision alternatives.

Table 3. Logical relations between architectural decision alternatives and their properties

Relation	Set(s)	Reflexive/ symmetric/ transitive	Cardinality (function?)	Other properties
<i>forces</i>	$A \times A$	no/no/yes	n:m (no)	Still a DAG w.r.t. $forces \cup <$ relations, which does not have to be connected (spawns a subgraph, which is not always a tree)
<i>isIncompatibleWith</i>	$A \times A$	no/yes/no	n:m (no)	Still a graph w.r.t. $isIncompatibleWith \cup <$ relations, which does not have to be connected (spawns a subgraph, which is not always a tree)
<i>isCompatibleWith</i>	$A \times A$	yes/yes/yes	n:m (no)	Default if no other relation exists between two alternatives (spawns a subgraph, which is not always a tree)

Our next two integrity constraints pertain to these three relations.

Integrity Constraint 3 A *forces* relation implies that an alternative in one issue $isIncompatibleWith$ all other alternatives in that decision:

$$\forall a_i, a_j, a_k, i_j < a_j, i_j < a_k, j \neq k: a_i \text{ forces } a_j \Rightarrow a_i \text{ isIncompatibleWith } a_k$$

Integrity Constraint 4 The *forces*, *isIncompatibleWith*, and *isCompatibleWith* relations between alternatives are mutually exclusive; one of them must exist. If nothing is defined, *isCompatibleWith* is the default.

$$\begin{aligned} \forall a_i, a_j: a_i \text{ forces } a_j \wedge a_i \text{ isIncompatibleWith } a_j &\equiv \text{false} \\ \forall a_i, a_j: a_i \text{ isIncompatibleWith } a_j \wedge a_i \text{ isCompatibleWith } a_j &\equiv \text{false} \\ \forall a_i, a_j: a_i \text{ forces } a_j \wedge a_i \text{ isCompatibleWith } a_j &\equiv \text{false} \\ \forall a_i, a_j: a_i \text{ forces } a_j \vee a_i \text{ isIncompatibleWith } a_j \vee a_i \text{ isCompatibleWith } a_j &\equiv \text{true} \end{aligned}$$

Rationale and example: We compare these relations with those defined in the ontology from Kruchten et al. in Section 6.

The *isIncompatibleWith* relation expresses that certain combinations of alternatives do not work with each other, for instance a NON-TRANSACTIONAL BACKEND service provider (not shown in Figure 5) can not be called from a service consumer that has

been decided to share transaction context with its provider (i.e., PAT decision to JOIN in Figure 5). A *forces* relation specifies that an alternative can only be combined with (only works with) one alternative in a different issue. For example, a conceptual alternative to share transaction context (PAT decision to JOIN) requires the technology-level Enterprise JavaBean (EJB) transaction attribute to be set to TX_MANDATORY.

In addition to the four formally defined integrity constraints, several heuristics can also be defined for an architectural decision model \mathbb{M} .

Definition 10 (Balanced Architectural Decision Model) *An architectural decision model \mathbb{M} is balanced if and only if the following informally defined heuristics regarding its structural properties hold:*

1. \mathbb{M} has at least two, but not more than five levels.
2. Topic groups do not contain more than nine other topic groups and twelve issues.
3. On all but the lowest level, there is at least one issue that has an outgoing refinement relation
4. On all but the highest level, there is at least one issue that has an incoming refinement relation.
5. The maximum path length to get from the initial issue to any issue via the contains relation $<$ and the maximum path length to get from the initial issue to any issue via refinedBy or decomposesInto relations is ten.

Rationale and example: Quality attributes such as usability and consumability for humans (e.g., knowledge engineers, software architects) justify these heuristics: An unbalanced model is difficult to maintain (for the knowledge engineer) and consume (for the software architect) due to the many elements per topic group and lengthy reasoning paths. According to studies in cognitive science and user interface design, three [13] to seven (plus/minus two) [14] entries on each level of a hierarchy are considered consumable. Good practices in object-oriented design give similar advice for inheritance trees [17]. Heuristic 1 adopts this advice; heuristic 2 and 5 are more tolerant due to experience we gained during RADM for SOA creation and tool implementation (see Section 6). Seven to nine architectural layers are defined in many reference architectures, and we often find ten or more components in each layer of a component-oriented architecture. If the topic group hierarchy resembles the architectural layering and logical decomposition into components, it must be able to deal with such numbers of topics groups and issues. Figure 5 shows a balanced architectural decision model.

4.3 Temporal Relations/Constraints and Decision Making Process Support

We add a relation to our model \mathbb{M} that facilitates the decision making process conducted by the software architect. Unlike previous definitions, this relation is not binary and defined between nodes of different types.

Definition 11 (triggers relation) *Let \mathbb{M} be an architectural decision model.*

Let a_i, a_j be decision alternatives in \mathbb{M} , let i_k be an issue in \mathbb{M} , and let t_l be a topic group in \mathbb{M} .

- $\text{triggers}(a_i, i_k, t_l)$ with $\neg(i_k < a_i)$ and $t_l < i_k$. Choosing an architectural decision alternative a_i can trigger an issue i_k and with this it triggers the topic group t_l which contains the issue. Indirectly, with the issue, all possible alternatives are triggered to direct the architect in the decision making process to the next recommended focus point, i.e., an issue where the next alternative should be selected. The relation adds additional directed edges to the model. The relation must not form any cycles when combined with $i_k < a_j$.

Table 4. Temporal relation in architectural decision models and its properties

Relation	Set(s)	Reflexive/ symmetric/ transitive	Cardinality (function?)	Other properties
<i>triggers</i>	$A \times I \times T$	n/a	n:m:1 (no)	Forms one or several DAGs, but not a tree.

Rationale and example: The *triggers* relation expresses a causal and therefore also temporal ordering during the decision making process. As we will see in Section 5 on dependency patterns, it is often combined with *refinedBy* or *decomposesInto* relations to form certain dependency patterns. Note the suggestive nature: It is permitted to resolve issues that have not been triggered (yet) and multiple triggers may exist per issue. It is possible that an alternative and an issue (and containing topic group) do not have any *triggers* relation. It would be far too restrictive for the architect to define a strictly enforced decision ordering based on these relations.

The *triggers* relation must satisfy the following integrity constraints:

Integrity Constraint 5 If an issue i_i is *refinedBy* or *decomposesInto* another issue i_j then any alternative in i_i *triggers* i_j :

$$\forall i_i, i_j, a_i, i_i < a_i: i_i \text{ refinedBy } i_j \vee i_i \text{ decomposesInto } i_j \Rightarrow a_i \text{ triggers } i_j$$

Integrity Constraint 6 The *forces* relation between alternatives *implies* a trigger relation:

$$\forall i_i, a_i, a_j: i_i < a_i \wedge i_j < a_j \wedge a_i \text{ forces } a_j \Rightarrow a_i \text{ triggers } i_j$$

In the next step, we define two more integrity constraints regarding the *triggers* relation. The logical implications caused by integrity constraints 5 and 6 allow us to define these solely on *trigger* relations (it is not required to include *refinedBy*, *decomposesInto*, and *forces* in the definitions):

Integrity Constraint 7 (Trigger Compatibility) Let a_i trigger i_j hold. Let $I(a_i)$ be the set of issues that can be reached from a_i following trigger relations and the contains relation $<$ within one tree \mathbb{T}_{km} starting with alternative a_i . Note that $I(a_i)$ can reach into other trees \mathbb{T}_m .⁹

⁹ $I(a_i)$ can be calculated like this: Initialize $I(a_i)$ with all issues triggered by a_i . Iterate: For any issue i added in the last iteration, follow the *triggers* relations originating in alternatives contained in i and add the target issues. Re-iterate if any issues were added in this iteration.

Then a_i must either have an `isCompatibleWith` relation with at least one alternative a_x or a `forces` relation with exactly one a_x for every $i_j \in I(a_i)$ and $i_j \prec a_x$:

$$\forall a_i, a_x \in A \forall i_j \in I(a_i): \\ i_j \prec a_x \Rightarrow a_i \text{ isCompatibleWith } a_x \vee a_i \text{ forces } a_x$$

Integrity Constraint 8 (Top-Down Progression) Let $i_i \prec a_i$ and a_i triggers i_j . i_j must then reside on a lower level than i_i or, if i_i and i_j reside on the same level, i_j must be greater than i_i according to \prec .

Rationale and example: Certain combinations of `triggers`, `isIncompatibleWith`, and `forces` relations should not occur. For example, an alternative must not trigger the issue in which it is contained (\prec relation). Less obvious consistency problems can occur when chaining more issues and alternatives together.

While a top-down approach to architecture design is taken in many methods, it can not always be applied in practice. When modernizing enterprise applications, many technology- and vendor asset-level decisions have already been made prior to project start (e.g., those pertaining to legacy systems). When procuring a software package, it comes with a certain approach to interface, transaction, and session management design. When deciding for a certain application server strategically, a vendor asset level decision is upgraded to the executive level. An architectural decision model for such a setting does not satisfy integrity constraint 8 (top-down progression). Hence, integrity constraint 8 is not always met in practice.

Definition 12 (Valid and Strictly Valid Architectural Decision Model) An architectural decision model \mathcal{M} is called valid if integrity constraints 1 to 7 hold. If \mathcal{M} is valid and integrity constraint 8 also holds, \mathcal{M} is called strictly valid.

Rationale and example: The transaction management example in Figure 5 meets all constraints. Therefore, it is a strictly valid architectural decision model.

Figure 6 illustrates several modeling errors. The model is not balanced due to the cyclic refines relations (i_1, i_2, i_3) , violating Definition 10 (part 3). It is not valid, either: a_{21} forces a_{13} and can therefore not be *compatible* with a_{12} (integrity constraint 3). Alternatives a_{12} and a_{21} can either be *compatible* or *incompatible* but not both (integrity constraint 4). i_2 refinedBy i_1 violates integrity constraint 8 due to the `triggers` relation implied by integrity constraint 5. a_{21} forces a_{13} implies a_{21} triggers i_1 (integrity constraint 6), but the implied `triggers` are not present in the model. a_{32} triggers i_4 , but there is no compatible alternative (as required by integrity constraint 7). a_{32} triggers i_2 which resides in a higher level (violating integrity constraint 8).

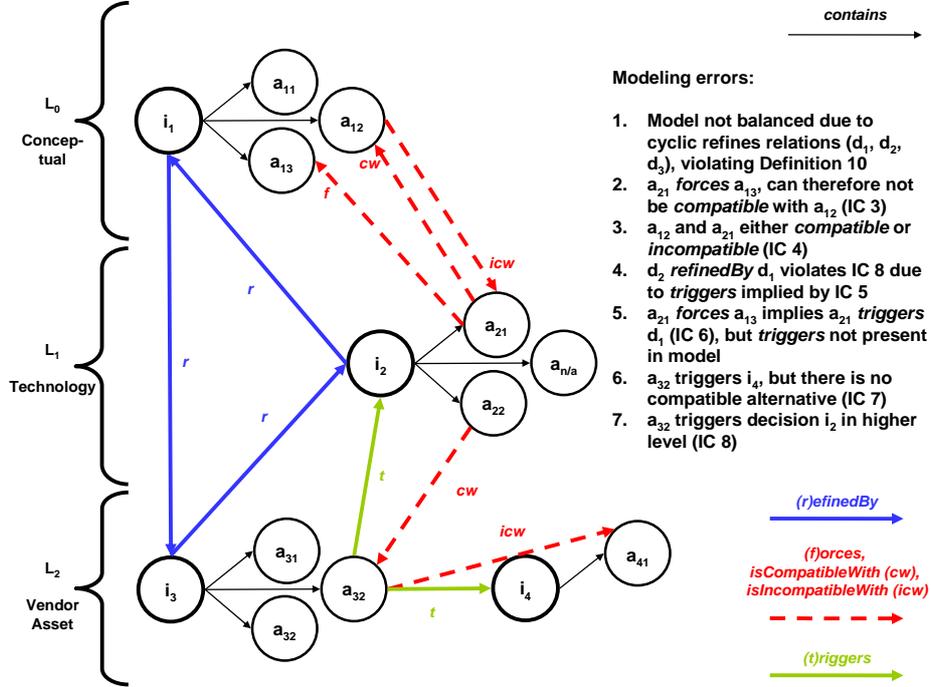


Fig. 6. Sample decision model violating integrity constraints.

Decision making process support. So far, we focused on modeling reusable architectural decision knowledge. We can now define how architectural decision models can be traversed on projects: We first define where to begin with the decision making and formalize ADOutcomes, which we then classify by their processing status determined by triggers relations.

Definition 13 (Entry Points, EP) The architectural decision Entry Points (EP) are the set of architectural decision issues in a architectural decision model \mathcal{M} that do not have any incoming triggers relations:

$$EP = \{ i \in I \mid \nexists a \in A: (a \text{ triggers } i) \}$$

Rationale and example: Entry points are a natural starting point for architecture design activities in a given project or project phase. There can be multiple ones. In Figure 5, the INVOCATION TRANSACTIONALITY PATTERN decision is the only entry point, which is marked as such. Note that the triggers can be implied by decomposesInto or refinedBy relations.

As we motivated in the example in Section 3, certain issues may have to be resolved multiple times, e.g., if the architecture applies a pattern such as “business process” or “service” multiple times. Each outcome captures a single decision made to resolve an issue. Hence, the UML metamodel from Section 3 specifies the dependency relation from ADIssue to ADOutcome to be 1 : n. In the formalization of the metamodel, this multiplicity is not defined yet. We add this support now:

Definition 14 (Outcome Instances, O, Open and Resolved Instances) Let O be a set of outcome instances $O = (name, candidateAlternatives, status) / name \in Strings, candidateAlternatives \subseteq A, status \in \{open, implied, resolved\}$ in a valid architectural decision model \mathcal{M} where $name$ indicates which element in the architecture is affected by the outcome, $candidateAlternatives$ is the subset of the alternatives contained in the issue to be considered for this outcome, and $status$ is a marking that is *open* initially and becomes *resolved* to indicate that an alternative has eventually been chosen by the architect. If $status$ is *open*, the outcome instance is called open outcome instance; if it is *resolved*, it is called resolved outcome instance. An *implied* status indicates that the decision can be concluded due to logical relations with outcome instances that have been resolved elsewhere.

Rationale: Outcome instances can be created to represent multiple occurrences of an issue in a project (recall the business process example in Section 3); their introduction models the transition from capturing reusable architectural knowledge (issues, alternatives) to the project-specific usage of this knowledge. Outcome instance names can either reference textual element identifiers in design models (e.g., business processes and Web services in SOA design) or integrate elaborate decision scoping concepts such as those described by Jansen and Bosch [5]. Outcome instances preserve and extend the tree structure of ADMs:

Definition 15 (hasOutcome relation \prec_o) Let $\prec_o \subseteq I \times O$ be a hasOutcome relation defined between issues and outcome instances. The cardinality of the relation is 1:n.

Rationale: An issue can be resolved by multiple outcome instances, but each outcome instance resolves exactly one issue and chooses exactly one alternative. Outcome instances are created on a project; the $candidateAlternatives$ attribute is set to all alternatives contained in the issue initially. During decision making, alternatives that cannot be chosen or are rejected (for whatever reason) are pruned from the $candidateAlternatives$ attribute until zero or one alternatives remain which means that the outcome instance can be implied or resolved by the architect.

Definition 16 (Open and Resolved Issue) An open issue is an issue which has a hasOutcome relation with at least one open outcome instance. A resolved issue a.k.a. decision made is an issue whose outcome instances are all resolved.

Rationale and example: Figure 7 adds six outcome instances to two of the issues from the previous example. Two of the outcome instances and both issues are open.

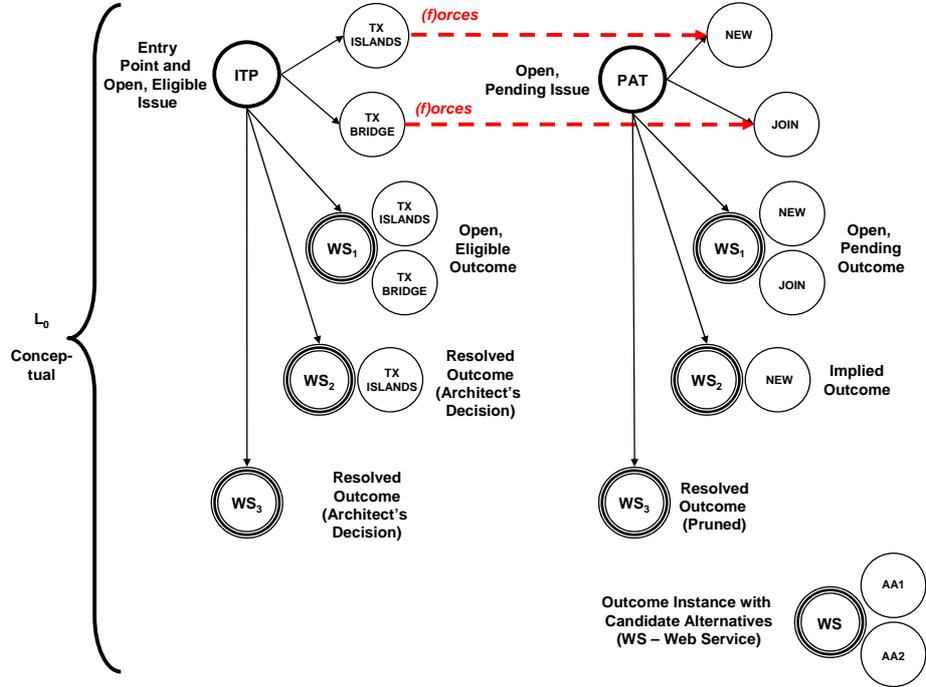


Fig. 7. Eligible and pending outcome instances in transaction management example.

Figure 7 classifies issues and outcome instances not only into open and resolved ones, but even further into eligible and pending ones:

Definition 17 (Eligible and Pending Outcome Instance) Let o_i and o_j be open outcome instances in an architectural decision model that have the same name (i.e., they refer to the same architecture element). Let i_i be the issue containing o_i . Let a_j be the alternative contained in issue i_j that contains the open outcome instance o_j . We call o_i an eligible outcome instance if there is no trigger relation from a_j to i_i . We call o_i a pending outcome instance if there is a trigger relation from a_j to i_i .

Rationale and example: All outcome instances are either eligible or pending. Eligible outcome instances can be resolved in the next decision making step, while pending ones have to wait until the ones they depend on have been made. Note that outcome instances can be eligible or pending because of triggers relations implied by *refinedBy* and *decomposesInto* relations.

Definition 18 (Eligible and Pending Issue) An open issue is eligible if it contains at least one eligible outcome instance. An open issue is pending if all contained outcome instances are pending.

Rationale: All open issues are either eligible or pending. Our approach is in line with the reasoning of Ran and Kuusela, who propose to start from issues that least likely have to be reverted during the decision making due to their dependencies. Many other classification principles exist, which are not included in our model yet (e.g., urgency of stakeholder request, related development effort, and technical risk).

In some cases, alternatives no longer have to be considered because of resolved outcome instances whose alternatives have *isIncompatibleWith* relations with each other. We now define three *production rules* to define such reasoning:

Production Rule 1 (Alternative Pruning) *If two alternatives have an *isIncompatibleWith* relation with each other and one of them is chosen during the decision making process in a resolved outcome instance, then it prunes the other alternative from the candidateAlternatives attribute in all outcome instances with the same name in which the other alternative appears:*

$$\begin{aligned} & \forall o_i, o_j \in O, a_i, a_j \in A: \\ & o_i.\text{candidateAlternatives} \equiv \{a_i\} \wedge o_i.\text{status} \equiv \text{resolved} \\ & \quad \wedge o_i.\text{name} \equiv o_j.\text{name} \\ & \quad \wedge a_i \text{ isIncompatibleWith } a_j \\ \Rightarrow & o_j.\text{candidateAlternatives} = o_j.\text{candidateAlternatives} \setminus \{a_j\} \end{aligned}$$

Rationale and example: For example, when a certain integration technology such as RESTFUL INTEGRATION is decided for, follow-up issues such as URI DESIGN and HIGH OR LOW REST are triggered, while all WSDL PORT TYPE alternatives become irrelevant and can be pruned from the candidate alternatives.

In some cases, the alternative to be chosen can even be implied:

Production Rule 2 (Outcome Implication) *If an alternative a_i appears in the candidateAlternatives of a resolved outcome instance, and it forces another alternative a_j , then all outcome instances with the same name that have a_j in their candidateAlternative set must chose a_j (i.e., all other alternatives can be pruned):*

$$\begin{aligned} & \forall o_i, o_j \in O, a_i, a_j \in A: \\ & o_i.\text{candidateAlternatives} \equiv \{a_i\} \wedge o_i.\text{status} \equiv \text{resolved} \\ & \quad \wedge o_i.\text{name} \equiv o_j.\text{name} \\ & \quad \wedge a_i \text{ forces } a_j \wedge a_j \in o_j.\text{candidateAlternatives} \\ \Rightarrow & o_j.\text{candidateAlternatives} = \{a_j\} \wedge o_j.\text{status} \equiv \text{implied} \end{aligned}$$

Rationale and example: In Figure 7, the PAT outcomes can be implied from the ones for INVOCATION TRANSACTIONALITY PATTERN (as a *forces* relation is present). This has happened for the outcome instance WS₂.

Integrity Constraint 9 *Only alternatives that do not have an *isIncompatibleWith* relation can be chosen within outcome instances that have the same name (i.e., either an *isCompatibleWith* or a *forces* relation must exist between the chosen alternatives):*

$$\begin{aligned} & \forall o_i, o_j \in O, a_i, a_j \in A: \\ & o_i.\text{candidateAlternatives} \equiv \{a_i\} \wedge o_i.\text{status} \equiv \text{resolved} \\ & \wedge o_j.\text{candidateAlternatives} \equiv \{a_j\} \wedge o_j.\text{status} \equiv \text{resolved} \\ & \quad \wedge o_i.\text{name} \equiv o_j.\text{name} \\ \Rightarrow & (a_i \text{ isCompatibleWith } a_j \vee a_i \text{ forces } a_j) \end{aligned}$$

Rationale and example: The six outcome instances in Figure 7 adhere to this IC.

Definition 19 (Implied and Pruned Outcome Instances) *An implied outcome is an outcome instance with all but one alternative pruned from the candidateAlternatives*

due to PR1 or PR2. A pruned outcome is an outcome instance with an empty set of candidateAlternatives, i.e., all alternatives have been pruned or removed manually.

Rationale and example: Figure 7 shows an implied and a pruned outcome instance. The existence of pruned outcomes merely expresses that the issue is not applicable for the architecture elements referred in its name (as classified by the scope attribute of the containing issue). It does not mean that a design is incomplete or erroneous, as successfully resolved outcome instances of the same name may be contained in other issues. We do not model such dependencies between outcome instances here; this requires further extensions of the formalization (e.g., formalize viewpoints and define cross-cutting integrity constraints). Such extensions are subject to future work.

Production Rule 3 (Outcome Instance Status Update) *If an outcome instance is implied or pruned, its outcome status is set from open to implied:*

$$\begin{aligned} & \forall o_i \in O, a_i \in A: \\ & o_i.\text{candidateAlternatives} \equiv \{\} \vee o_i.\text{candidateAlternatives} \equiv \{a_i\} \\ & \Rightarrow o_j.\text{status} \equiv \text{implied} \end{aligned}$$

Rationale: The architect still has to confirm that the implication is technically sound; it might as well be necessary to backtrack and revise a related decision that has been made previously. Hence, PR3 sets the outcome instance to an intermediate state implied and not to resolved.

Definition 20 (Pruned Issue, Pruned Topic) *If all outcome instances of an issue are pruned outcome instances, the issue is called pruned issue; if a topic only contains pruned issues, it is called a pruned topic.*

We can verify whether additional decision making is still required.

Definition 21 (Decided, Correct Architectural Decision Model) *A valid decision model is called decided if all outcome instances are resolved and, in turn, no open issues exists. If integrity constraint 9 holds, the decided model is called correct.*

Rationale and example: When the decision making process completes, all decisions must have been made, i.e., neither eligible nor pending open issues exist.

With these definitions in place, the decision making process can be characterized as follows, showing mixed initiatives by the architect A and a decision support system S implementing the concepts defined in this section:

```
decide (in: strictly valid decision model,
       out: decided decision model)

[S: set initially eligible decisions to entry points]
While [decision model is not decided (Def. 21)]
  For [all eligible issues and outcome instances (Def. 17)]
    [A: Group issues/instances by scope/phase/role (Def. 2)]
    [A: Make decisions in each group]
    If [S: decision model violates IC 9]
      [A: Reset selected outcome instances to open]
      [A: Choose other alternatives]
      Continue (with If)
    Else
      [S: Prune alternatives (PR 1)]
      [S: Imply outcome instances (PR 2)]
      [S: Update outcome instance status (PR 3)]
```

```

    [A: Resolve/confirm implied outcome instances]
  End if
End for
[S: Calculate eligible outcome instances and issues]
End while

```

5 Dependency Patterns

In this section, we generalize the SOA decision modeling examples introduced so far into broadly applicable *dependency patterns*. The patterns combine certain decision types introduced in Section 3 with certain instances of *refinedBy*, *isIncompatibleWith*, *forces*, and *triggers* relations defined in Section 4.

Figure 8 introduces a second decision modeling example, the design of an integration architecture starting with the classical BROKER pattern. T

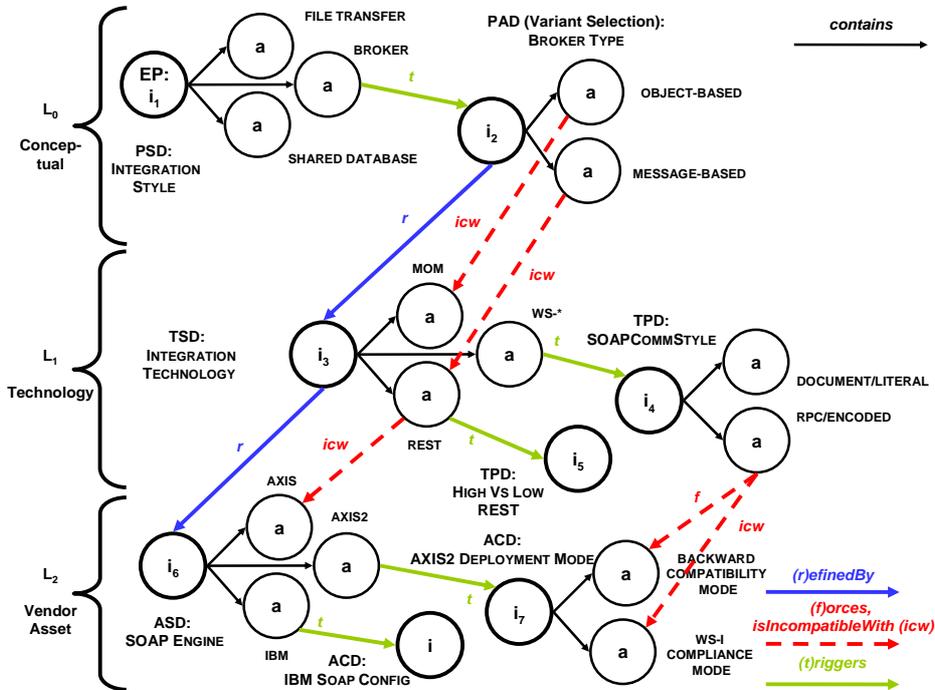


Fig. 8. Refinement and decomposition of pattern adoption decision about integration broker.

The model is a strictly valid architectural decision model adhering to all integrity constraints defined in Section 4. The same three levels as in the previous example shown in Figure 5 are defined. Several instances of the decision types introduced in Table 1 are present. The architectural PSD about an INTEGRATION STYLE in the conceptual level is the only entry point; one of its alternatives has an outgoing *triggers* relation with an architectural PAD regarding a pattern variant on the same level (BROKER TYPE). The pattern variants are modeled as alternatives of the architectural PAD. They constrain the possible choices for the technology decisions

(TSDs and TPDs) on the technology level. Here, the architectural PAD is *refinedBy* a TSD INTEGRATION TECHNOLOGY. Its WS-* alternative triggers one TPD, SOAP COMM STYLE. The REST alternative triggers another TPD HIGH VS LOW REST. If the INTEGRATION TECHNOLOGY is WS-* and not REST, there is no need to decide for a certain URI design style, which is the scope of the HIGH VS LOW REST decision.¹⁰

In Figure 8, the relations between the technology level and the vendor asset level resemble those between the conceptual level and the technology level. The TSD INTEGRATION TECHNOLOGY is *refinedBy* a vendor ASD SOAP ENGINE, which *triggers* a vendor ACD AXIS2 DEPLOYMENT MODE; the rationale is that different SOAP engines require different proprietary ACDs. These are the first two examples of a recurring *dependency pattern*. The *refinedBy* and the *forces* correspondences between PAT and the contained JOIN alternative and IAT and PARTICIPATES in Figure 5 in Section 3 can also be seen as instances of this pattern. A fourth instance of this pattern can be observed between CT and TRANSPORT QOS, also in Figure 5. In this case, the originating decision resides on the conceptual level and, unlike in the other pattern instances, the destination resides on the vendor asset level.

Figure 9 generalizes these examples of cross-level dependencies, commonly occurring between certain types of decisions, into two dependency patterns, TECHNOLOGY LIMITATION and PRODUCT LIMITATION:

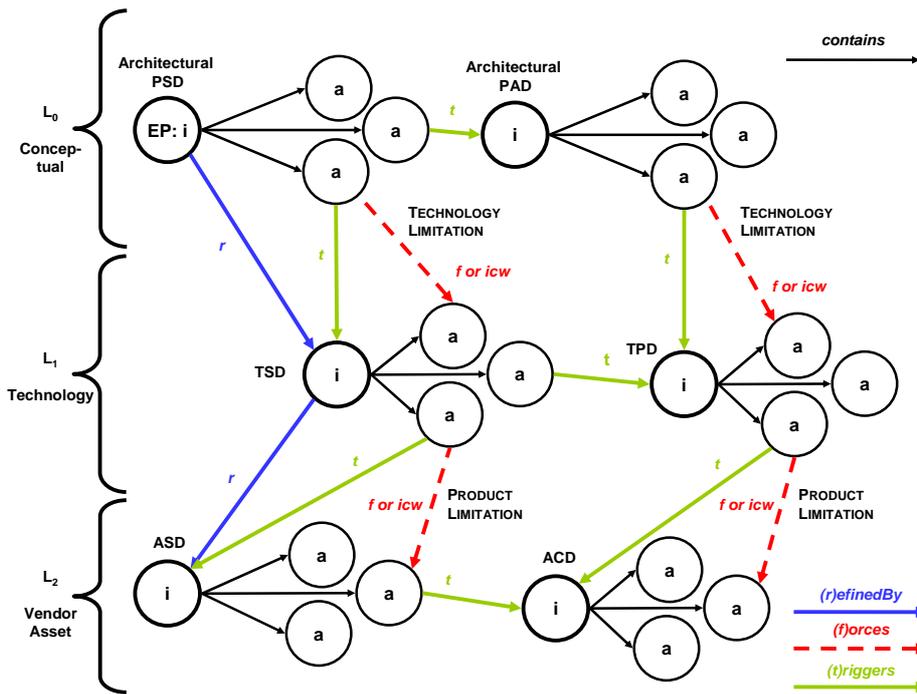


Fig. 9. Decision making patterns (top down): TECHNOLOGY LIMITATION, PRODUCT LIMITATION.

¹⁰ Not all relations that exist in the real model are shown in the figure and explained in the text (in the interest of readability).

6 Model Analysis, Practical Use, and Implementation

As already mentioned in the introduction and in Section 2, the formalization presented in this paper has its origins in an industrial knowledge engineering project we have been conducting since January 2006, *SOA Decision Modeling (SOAD)*. SOAD has three project objectives and types of results:

1. Defining the fundamental concepts of a decision-centric architecture design *method*. Sections 3 and 4 of this paper contribute a domain metamodel and a formalization of decision dependencies and integrity constraints to this method. The application of the method to enterprise application design and the relationship with pattern languages is addressed in [26][27].
2. Providing reusable decision *content* (knowledge) for SOA construction projects. Selected decisions from this RADM for SOA already served as examples in this paper (Sections 3 to 5). More decisions are featured in other publications [15][24][26].
3. Demonstrating how the decision modeling concepts can be implemented and how the decision content can be managed collaboratively with the help of a *tool*. Architectural Decision Knowledge Wiki, made publicly available in March 2008, serves this purpose [19].

We validate our research results by analysis, implementation and experiment, as well as industrial case studies involving action research. To analyze the maturity of the domain metamodel, Section 6.1 compares our dependency modeling with an existing taxonomy. The SOA content and the tool support are two more means of validation for the SOAD concepts, particularly the ones presented in this paper. We give an overview of the content and tool validation results in Sections 6.2 and 6.3.

6.1 Comparison with a State-of-the-Art Taxonomy

Table 5 compares the dependency types from [12] with those from Section 4.

Table 5. Dependencies defined by Kruchten et al. and their representation in our formal model

Relation type	Corresponding relation in our model	Comparison and assessment
Constrains	<i>forces, isCompatibleWith</i> plus integrity constraints	Our approach as defined in Sections 3 and 4 is slightly more elaborate
Forbids	<i>isIncompatibleWith</i> , pruning	Our approach separates logical and temporal aspects
Enables	<i>triggers</i>	Same concept, but two entities appear in our approach (issue , alternative)
Subsumes	<i>refinedBy, decomposesInto</i> plus integrity constraints	Our approach is slightly more elaborate, using the level concept
ConflictsWith	<i>isIncompatibleWith</i> , pruning	Same concept, but an additional entity is used (alternative)
Overrides	Compares to concepts of outcome instances	Can be expressed with ADOOutcome concept from UML metamodel
Comprises	<i>decomposesInto</i>	Same concept, inverse direction
IsAnAlternativeTo	<i>contains</i> relation, alternatives with same parent decision (siblings)	Not between alternatives, but between ADIssue and ADAAlternative instances

		in our approach (same expressivity)
IsBoundTo	ADTopicGroup node, decision scoping concept [25]	Approaches have similar modeling capabilities, but use different entities
IsRelatedTo	<i>influences</i>	Same expressivity

A major difference is that Kruchten et al. define binary relations over a single entity, namely the decision, whereas our UML metamodel defines five classes: ADLevel, ADTopicGroup, ADIssue, ADAAlternative, and ADOOutcome. As the table shows, the semantics of the various dependency relations, however, is very similar. Our model is formally defined. For five of Kruchten's relations, we provide more elaborate modeling concepts due to the explicit representation of alternatives.

6.2 Practical Use of Modeling Concepts: Reusable ADM for SOA

We applied the UML metamodel from Section 3 and the formal modeling concepts from Section 4 to enterprise application development and SOA design to produce the second result of the SOAD project, content. Our *Reusable Decision Model (RADM) for SOA* is a balanced architectural decision model with four levels which is strictly valid. Its initial content originated from several large-scale SOA development projects conducted from 2001 to 2005 [1]. Since then, the content was extended and refactored several times; knowledge from a practitioner community was incorporated (30+ projects, 200 decisions). The metamodel remained stable. Figure 11 outlines the structure of the RADM:

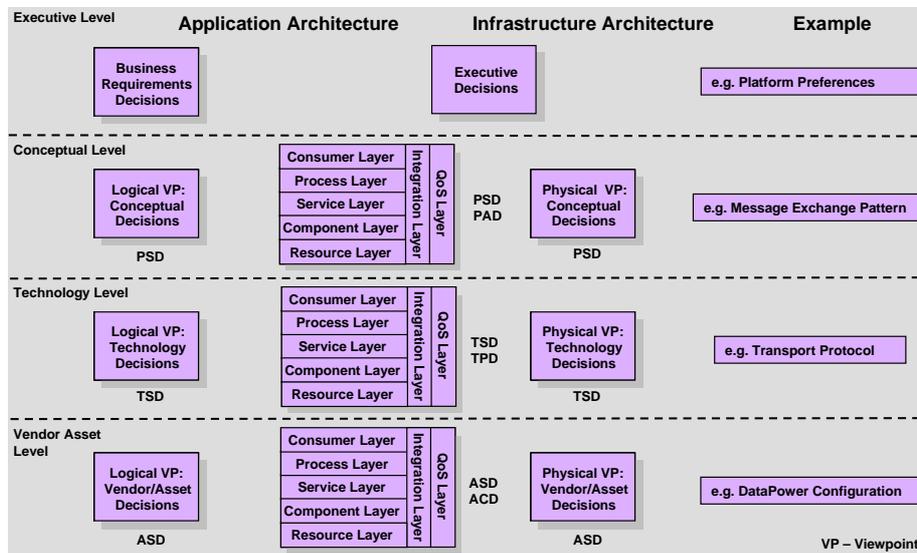


Fig. 11. Layers and levels in RADM for SOA.

The four levels were introduced in Section 3. Each box represents one topic group. The same top-level topic groups are defined on the conceptual, the technology, and the vendor asset level: They represent seven logical layers (consumer, process, service, component, resource, integration, QoS layer). For instance, one of these topic

groups contains issues about the consumer layer on the conceptual level. Two topic groups on each level contain issues pertaining to the logical and physical viewpoint that can not be assigned to any layer. All decision types from Table 1 in Section 3 are used. Various relations as defined in Section 4 are present. All integrity constraints defined in Section 4 are met, including trigger compatibility and top-down progression. A GO NO GO DECISION serves as a single global entry point. The model can be tailored and irrelevant parts removed, e.g., if only decisions dealing with layer 5 processes (workflows) are of interest in a particular project context. After such tailoring step, new entry points become available, residing in the conceptual logical viewpoint topic groups. About a dozen subject area keywords are defined and expressed as topic tags (which is an ADIssue attribute according to Definition 2), e.g., session management, transaction management, and error handling.

At present, the RADM for SOA consists of 86 topic groups and 389 issues with ~2000 alternatives. The knowledge base is still growing, now at a slower pace than in the beginning of the project. While this growth could continue forever (at least in theory), we plan to freeze the knowledge engineering once the 500 most relevant issues have been compiled. The knowledge base will still have to be reviewed periodically to ensure that the contained information remains up to date. Issues and alternatives will become obsolete as technology evolves; new ones will be required. The knowledge engineer can utilize the dependency relations, integrity constraints, and structural heuristics defined in this paper during this maintenance process.

6.3 Tool Implementation: Architectural Decision Knowledge Wiki

Architectural Decision Knowledge Wiki is a model-based collaboration system that implements the domain metamodel defined in Section 3. The central concept is the architectural decision model from Definition 7 in Section 4. The levels are freely configurable; users are not obliged to stick to the conceptual, technology, vendor asset level structure used in this paper and in the RADM for SOA. These levels, however, have proven to be appropriate for structuring the SOA content.

The main content structuring principle is the level and topic group hierarchy. At present, Architectural Decision Knowledge Wiki supports about 50 use cases, providing decision modeling functionality in the following areas:

- Import and export of decision content.
- Create, read, update, and delete operations on ADTopicGroups, ADIssues, ADAlternatives, and ADOutcome instances.
- Decision lifecycle management and community involvement.
- Relationship editor.
- Search and filter by role, phase, and scope attributes and by topic tag.
- Report generation.

All model attributes can queried, e.g., the status, role, and phase attributes when looking for all open decisions an application architect has to make in the solution outline (a.k.a. inception) phase of an enterprise application development project. The integrity constraint checks and heuristics defined for balanced architectural decision

models from Definition 10 are implemented in an advanced prototype that is not yet publicly available.

Figure 12 shows a screen capture of the page displaying ADIssue and ADAAlternative instances:

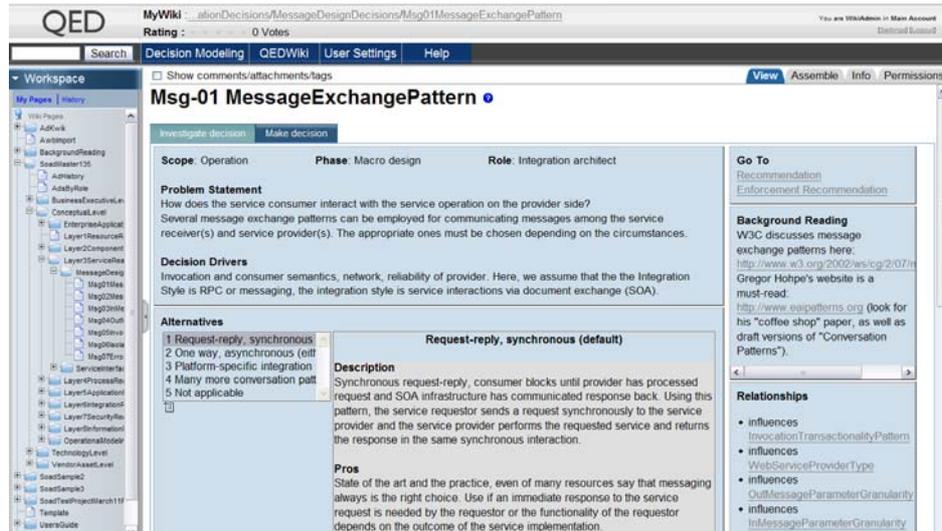


Fig. 12. Screen capture of Architectural Decision Knowledge Wiki.

Architectural Decision Knowledge Wiki is available on IBM alphaWorks [19]; an earlier version of it is described in detail in a separate publication [18]. The tool has already been used in several industrial projects and training classes. More than 150 users are registered in a company-internal hosted instance. More than 400 interested parties downloaded the tool from IBM alphaWorks.

Many change cases have already been identified based on feedback from early adopters. For instance, the containment-oriented view shown on the left in Figure 12 was not well received. Therefore, we designed an additional *AD Status Overview* view. This view makes use of the classification of decisions into entry points, eligible, pending, and implied issues as introduced in definitions 13 to 20. Integration with other tools used by architects, for example UML modeling environments, requirements engineering tools, and development team collaboration platforms, was also requested as a future extension.

7 Conclusion and Outlook

In this paper, we presented a formal model for capturing and reusing architectural decision knowledge. Our approach extends existing proposals for retrospective architectural decision capturing with a formal definition of architectural decision models and modeling concepts for collaboration and reuse, e.g., levels, dependency relations, integrity constraints, and production rules. We used this model to capture 389 SOA issues. Decision types such as executive decisions, pattern selection and

adoption decisions, technology selection and profiling decisions, asset selection and configuration decisions appear in this model. Selected decisions from this SOA decision model served us as examples.

The decision types introduced in Section 3, the relations defined in Section 4 and the dependency patterns from Section 5 serve several purposes: First, they can help knowledge engineers and software architects to detect design flaws (in reusable assets, on individual development and integration projects). Furthermore, they have educational character for consumers of architectural knowledge. Decision identification, making, and enforcement tools can be built that guide decision makers through their activities and verify integrity constraints along the way. Pruning can be used to cut off alternatives and entire sub trees when a decision is made. This simplifies the rather complex task of managing a complex decision model.

Future work concerns formalizing additional characteristics of tree-based architectural decision models and the relationship between decision models and other model types used to document the various views on software architecture such as Kruchten's 4+1 views [10]. The design fragments from Jansen and Bosch and the SPEM integration from de Boer et al. can be leveraged to do so. Additional constraints on various relations can be expressed. Finally, integrating SOAD with natural controlled language such as Attempto Controlled English (ACE) [7] is another promising area of future research: If SOAD decision drivers and related best practices recommendations are articulated in a natural controlled language such as ACE, a reasoning engine can analyze them and suggest certain alternatives to the architect.

We envision several advanced usage scenarios for the concepts presented in this paper. Project managers can use decision models for planning purposes. Work breakdown structures and effort estimation reports can be created, as open decisions correspond to required activities. Health checking is another application area: If there are many, frequent changes, or many questions are still unresolved in late project phases, the project is likely to be troubled. Product selection decisions define which software licenses are required, and on which hardware nodes the required software has to be installed. Moreover, the outcome of product-specific asset configuration decisions can serve as input to software configuration management. The model can also serve enterprise architects; they can maintain a company-specific instance of the decision model, consisting of a subset of issues and alternatives. Such an approach authorizes solution architects on projects to make decisions ("freedom of choice") without sacrificing architectural integrity ("freedom from choice"). Finally, the reusable decision model for SOA can be used as a supplemental design method for SOA construction which complements and details existing service modeling methods.

References

- [1] Abrams, S., Bloom, B., Keyser, P., Kimelman, D., Nelson, E., Neuberger, W., Roth, T., Simmonds, I., Tang, S., and Vlissides, J. Architectural thinking and modeling with the architects' workbench. *IBM Syst. J.* 45, 3 (Jul. 2006), pp. 481-500.
- [2] Akerman, A. and Tyree, J. Using ontology to support development of software architectures. *IBM Syst. J.* 45, 4 (Oct. 2006), pp. 813-825.

- [3] Bass, L.; Clements, P.; Kazman, R.: Software Architecture in Practice, Second Edition, Addison Wesley, 2003
- [4] de Boer R.C., Farenhorst, R., Lago P., van Vliet H., Clerc V., and Jansen A. Architectural Knowledge: Getting to the Core. In Third International Conference on Quality of Software-Architectures (QoSA), (Jul. 2007), pp. 197-214
- [5] Jansen, A. and Bosch, J. Software Architecture as a Set of Architectural Design Decisions. In Proceedings of the 5th Working IEEE/IFIP Conference on Software Architecture (Nov. 2005). IEEE Computer Society, Washington, DC, pp. 109-120
- [6] Fowler M., Patterns of Enterprise Application Architecture, Addison Wesley, 2003
- [7] Fuchs, N. E. and Schwitter, R., Attempto Controlled English (ACE), in: Proceedings of CLAW 96, First International Workshop on Controlled Language Applications, University of Leuven, Belgium, (March 1996), pp. 124-136
- [8] Harrison, N. B., Avgeriou, P., and Zdun, U. 2007. Using Patterns to Capture Architectural Decisions. IEEE Softw. 24, 4 (July 2007), 38-45
- [9] Hohpe G., Woolf, B., Enterprise Integration Patterns, Addison Wesley, 2004.
- [10] Kruchten P., The 4+1 View Model of Architecture, IEEE Software, vol. 12, no. 6, Nov. 1995, pp. 42-50
- [11] Kruchten P., The Rational Unified Process: An Introduction, Addison-Wesley, 2003
- [12] Kruchten P., Lago P., van Vliet H, Building up and reasoning about architectural knowledge. In: Hofmeister, C. (Ed.), Proceedings of Second International Conference on the Quality of Software Architectures (QoSA 2006), Springer LNCS 4214, 2006, pp. 43-58
- [13] LeCompte, D., Seven, plus or minus two, is too much to bear: Three (or fewer) is the real magic number, Proceedings of the Human Factors and Ergonomics Society, (1999), pp. 289-292
- [14] Miller, G.A., The magical number seven, plus or minus two: Some limits on our capacity for processing information, The Psychological Review, (1956), pp. 81-97
- [15] Pautasso C., Zimmermann O., Leymann F., RESTful Web Services vs. Big Web Services: Making the Right Architectural Decision. In: W.-Y. Ma, A. Tomkins, X. Zhang (eds.): Proc. of WWW 2008, ACM Press (2008), pp. 805-814
- [16] Ran A., Kuusela J., Design Decision Trees, in 8th International Workshop on Software Specification and Design, (1996), pp. 172-175.
- [17] Riel A. J., Object-Oriented Design Heuristics, 1st edition, Addison-Wesley Longman Publishing Co., Inc. Boston, MA, USA, 1996
- [18] Schuster N., Zimmermann O., Pautasso C., AD_{kwik}: Web 2.0 Collaboration System for Architectural Decision Engineering. In: Proceedings of the Nineteenth International Conference on Software Engineering & Knowledge Engineering (SEKE'2007), KSI, (2007), pp. 255-260
- [19] Schuster N., Zimmermann O., Architectural Decision Knowledge Wiki, IBM alphaWorks, March 2008, <http://www.alphaworks.ibm.com/tech/adkwik>
- [20] Tang, A., Babar, M. A., Gorton, I., and Han, J. A survey of architecture design rationale. J. Syst. Softw. 79, 12 (Dec. 2006), pp. 1792-1804
- [21] Tyree, J. and Akerman, A. Architecture Decisions: Demystifying Architecture. IEEE Softw. 22, 2 (Mar. 2005), pp. 19-27
- [22] Web Services Description Language (WSDL) 1.1, <http://www.w3.org/TR/2001/NOTE-wsdl-20010315>
- [23] Zdun U., Dustdar S., Model-Driven and Pattern-Based Integration of Process-Driven SOA Models, <http://drops.dagstuhl.de/opus/volltexte/2006/820>
- [24] Zimmermann, O., Doubrovski, V., Grundler, J., and Hogg, K. Service-oriented architecture and business process choreography in an order management scenario: rationale, concepts, lessons learned. In Companion To the 20th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (San

- Diego, CA, USA, October 16 - 20, 2005). OOPSLA '05. ACM, New York, NY, pp. 301-312.
- [25] Zimmermann O., Gschwind T., Küster, J., Leymann F., Schuster N., Reusable Architectural Decision Models for Enterprise Application Development. In: Overhage S., Szyperski C. (eds.), QOSA 2007. LNCS, Springer, Heidelberg (2007), pp. 15-32
 - [26] Zimmermann, O., Milinski, S., Craes, M., and Oellermann, F. Second generation web services-oriented architecture in production in the finance industry. In Companion to the 19th Annual ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages, and Applications (Vancouver, BC, CANADA, October 24 - 28, 2004). OOPSLA '04. ACM, New York, NY, pp. 283-289
 - [27] Zimmermann, O., Zdun, U., Gschwind, T., and Leymann, F. Combining Pattern Languages and Reusable Architectural Decision Models into a Comprehensive and Comprehensible Design Method. In Proceedings of the Seventh Working IEEE/IFIP Conference on Software Architecture (WICSA 2008) - Volume 00 (February 18 - 21, 2008). WICSA. IEEE Computer Society, Washington, DC, pp. 157-166