

# Automated Resolution of Connector Architectures Using Constraint Solving (ARCAS method)

Jaroslav Keznikl<sup>1,2</sup>, Tomáš Bureš<sup>1,2</sup>, František Plášil<sup>1,2</sup>, Petr Hnětynka<sup>1</sup>

Technical report No. 2012/1, February 2012  
Version 1.0, February 2012

<sup>1</sup> Distributed Systems Research Group, Department of Software Engineering  
Faculty of Mathematics and Physics, Charles University  
Malostranské nám. 25, 118 00 Prague, Czech Republic  
phone +420-221914267, fax +420-221914323

<sup>2</sup> Institute of Computer Science, Academy of Sciences of the Czech Republic  
Pod Vodárenskou věží 2, 182 07 Prague, Czech Republic  
phone +420-266053831

# Automated Resolution of Connector Architectures Using Constraint Solving (ARCAS method) [DRAFT]

JAROSLAV KEZNIKL  
TOMÁŠ BUREŠ  
FRANTIŠEK PLÁŠIL  
PETR HNĚTYNKA

Charles University in Prague, Faculty of Mathematics and Physics  
Malostranske namesti 25, 118 00 Prague 1, Czech Republic  
{keznikl, bures, plasil, hnetynka}@d3s.mff.cuni.cz  
<http://d3s.mff.cuni.cz>

**Abstract:** In current software systems, connectors play an important role by encapsulating the communication and coordination logic. Since they share common patterns (elements) depending on characteristics of the connections, the elements can be predefined and reused. A method of connector implementation based on a composition of predefined elements naturally comprises two steps – resolution of the connector architecture, and creation of the actual connector code based on the architecture. However, manual resolution of connector architecture is not feasible due to the number of factors to be considered. Thus, the challenge is to come up with an automated method, able to address all the important factors. In this paper, we present a method for automated resolution of connector architectures based on constraint solving techniques (ARCAS). We exploit the Alloy modeling language for defining a connector theory, reflecting both the predefined parts and the important resolution factors, and employ a constraint solver to find a suitable connector architecture as a model of the theory.

**Keywords:** Software Architecture, Software Connectors, Constraint Solving, Middleware-based Connectors, Connector Theory, Alloy

## 1. INTRODUCTION

Proposed with the aim of supporting the separation of concerns, software connectors [MMP00, TMD10] are entities solely encapsulating communication and coordination among components. In particular, connectors ensure distribution of communicating components [BP04] while encapsulating middleware (*middleware-based* connectors), provide adaptation in order to achieve middleware-level [IBB11, NTER06] and application-level [SI10, CCP11, IST11] interoperability (*adaptors*), and ensure synchronization of component communication [BS07, IST11] (*coordinators*). In this paper, we focus particularly on the middleware-based connectors.

Although the introduction of middleware-based connectors provides benefits in terms of separation of concerns and abstraction of particular middleware, it does not necessarily simplify the code development effort, since, in principle, the middleware-related code is moved from components to connectors. In the component models that include connectors, e.g., [RCGT09, TMD10], the connector lifecycle differs from the component lifecycle: although partially-specified connectors are employed during the application design phase, fully-specified connectors emerge at the earliest in the component deployment phase – after decisions on application architecture and deployment have been made. Moreover, deployment of a particular component application may vary from time to time, so that several variants of a connector may be required. Advantageously, these variants typically share common patterns related to a particular communication style [BP04, TMD10, IBB11], middleware, and non-functional properties (NFPs); therefore, related parts of connectors can be predefined/designed in advance.

Middleware-based connectors can also emerge later, even after some of the components are already running. This is particularly true in the case of independently deployed components available as services (e.g., web services). In principle, the task of a newly emerging connector is to mediate a client component's communication with such a service while respecting the particular middleware employed by the service. In this sense, services are middleware-aware components. Similarly, the

emergence of such connectors can be desirable at runtime once architecture and deployment reconfiguration takes place (e.g., due to load-balancing).

In this context, the challenge is to find an automated method for synthesis of middleware-based connectors at deployment time/run time, i.e., synthesis of *emergent connectors* [ISJB09], in such a way that reuse of predefined connector parts is maximized. A related issue is to structure the predefined parts accordingly. Another challenge is to support NFPs in the actual connector synthesis and to structure the predefined parts in order to efficiently and flexibly capture variability of NFP requirements. The position of such automated connector synthesis is illustrated in Fig. 1.

### 1.1 Goal of the Paper

The goal of this paper is to respond to the challenges above by introducing the ARCAS method (Automated Resolution of Connector Architectures using a constraint-Solving technique). In general, ARCAS is based on an automated composition of connectors from predefined hierarchical *elements* [BP04, RCGT09]. It produces a description of a hierarchical composition of elements reflecting the connector design and deployment requirements imposed on the components being connected.

Technically, given a design specification of a connector, including NFPs and decision on component deployment, ARCAS produces a connector instance configuration (CIC), describing a particular composition of available elements to realize the connector. From a big-picture perspective, ARCAS is the first phase in a two-step connector generation process [B06]. In the second step, CIC is used as the input for the actual code generation process [MPBH11] yielding a deployable connector code.

The basic idea of ARCAS (Fig. 2) is to employ a constraint-solving technique for automated resolution of CIC. For this purpose, we employ the Alloy modeling

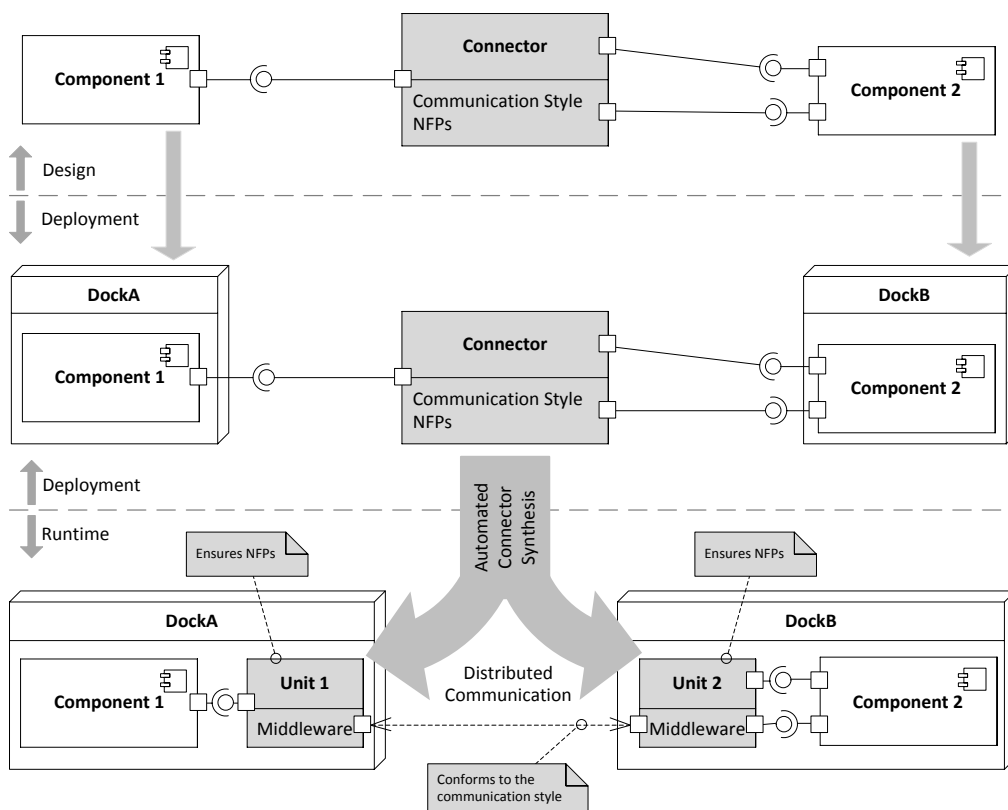


Fig. 1 Role of automated connector synthesis in connector lifecycle




Fig. 2 ARCAS overview

language [J02, J06] for capturing a *connector theory*, i.e., a logic theory playing the role of a constraint specification reflecting both the predefined elements, and the design and deployment requirements of the connector. Further, we employ the Alloy Analyzer as a constraint solver to find a model of the theory, representing a desired CIC. As an aside, the ARCAS method has applicability also in other areas dealing with configuration management of component-based applications and product-lines.

## 1.2 Structure of the Paper

The paper is structured as follows: Section 2 presents the basic concepts of middleware connectors and illustrates these with an example. Section 3 gives a brief overview of the whole ARCAS method. Section 4 describes both abstract and concrete syntax of a middleware-connector specification. Section 5 describes the construction of a connector theory in terms of predicate logic and relational calculus. Section 6 provides a brief introduction to the Alloy modeling language and describes ARCAS in terms of the Alloy modeling language. Section 7 surveys the related work, while Section 8 provides evaluation and discussion of the ARCAS method, and Section 9 concludes the paper while suggesting future work activities.

## 2. MIDDLEWARE CONNECTORS – BASIC CONCEPTS

In this section, we introduce the basic concepts of middleware connectors. A connector can be viewed from: (i) application-designer perspective (requirement and deployment views) and (ii) connector-designer perspective (design view). While (i) focuses on the high-level task and properties of a connector in a particular component application, (ii) focuses strictly on the connector design and implementation. Here, the concepts in (i) are generally agreed in the area of middleware-based connectors [CL02], whereas the concepts in (ii) stem from our experience with connector design and implementation [BP04] and are thus rather specific to ARCAS. We recall and illustrate all the concepts on a simple example – a fragment of a distributed component-based application [HKW08] featuring the components `CashDesk` and `Inventory` bound together by a single connector. This setting will be used as a running example throughout the text.

From the application-designer perspective, the requirements view of a connector (Fig. 3) focuses on describing the communication style and the required NFPs (henceforth referred to as *features*) of a component connection. The communication style defines *roles* – the connector’s endpoints for communication of components [CL02]; e.g., the procedure-call communication style defines the roles `Caller` and `Callee`. The components to be connected communicate using instances of the roles. Therefore, the requirements view of a connector has to comprise an association of the respective component interfaces with particular instances of connector’s roles; e.g., the `CashDesk`’s required interface is associated with an instance of the role `Caller` and the `Inventory`’s provided interface is associated with an instance of the role `Callee`. Finally, the requirements view of a connector defines the required features of the component connection; e.g., the requirement that all connector invocations have to be logged to a file and that the communication has to be secure. For the purpose of this




Fig. 3 Example of the application-designer perspective – requirements view

text, the requirements view of a connector is assumed to be described by a *requirements specification*.

At the deployment time, a connector has to reflect the distribution of the connected components according to the actual deployment decision. This is the focus of the deployment view of a connector (Fig. 4). In the example, the distributed environment consists of two component containers (*deployment docks*) A and B. Thus at the deployment time a connector is viewed as an assembly of distributed connector *units*. The key purpose of a unit is to refine the roles associated with a particular component, e.g., the connector defined in Fig. 4 is split into two units, each related to one of the CashDesk and Inventory components. In compliance with the desired component deployment, units have to conform to the *capabilities* of the selected deployment docks. Deployment dock capabilities are key properties of the execution environment, driving selections of middleware technology for remote communication and are based on the OMG D&C standard [OMG04]. For example, the capabilities of the dock “A” indicate the availability of Java virtual machine version JDK 1.4 and underlying Linux 2.4.28. Thus, the unit for the CashDesk component has to be able to run and communicate in such runtime environment. Information on the actual deployment decision and capabilities of all docks is given in a *deployment specification*.

From the connector-designer perspective, the design view of a connector focuses on describing the connector implementation by means of (possibly hierarchical) *elements*. An important idea is that the individual parts of connector implementation are designed in advance and reused (this is supported by the hierarchical structure of the connector design). Thus, all henceforth-introduced concepts facilitate composability and reuse. To allow design in advance, the particular application (i.e., the components and their deployment) is abstracted away by considering just communication style, features, and capabilities (thus, the communication style, features, and capabilities are the binding concepts of the application-designer and connector-designer perspectives). At the top level, the connector is described by its *distribution architecture* (Fig. 5), defining the potential units, i.e., the anticipated distribution of the connector, and the refinement of the roles by the units (e.g., the

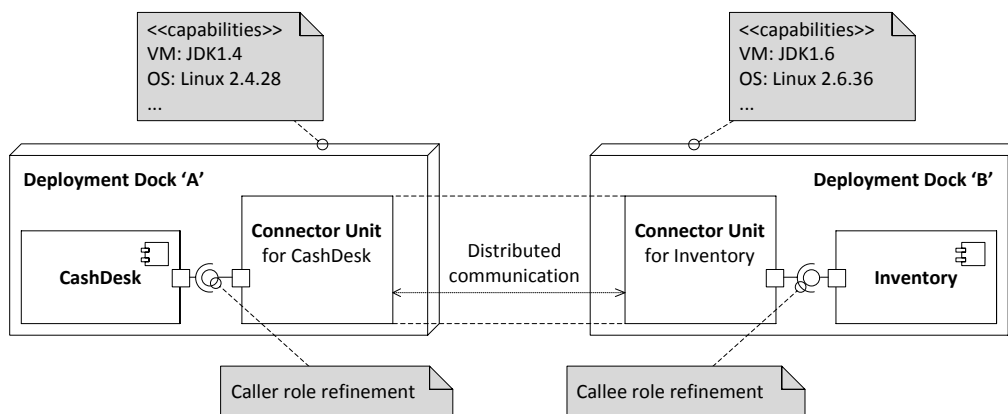


Fig. 4 Example of the application-designer perspective – deployment view

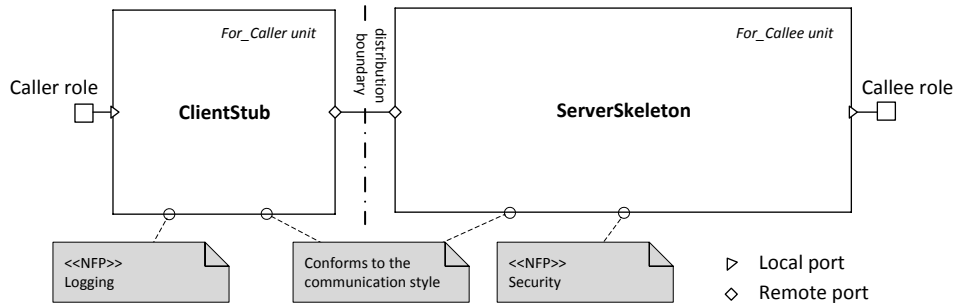


Fig. 5 Example of the distribution-architecture of a connector

units *For\_Callee* and *For\_Caller*). At this point, it also reflects the units as the top-level elements defined just by their type (*ClientStub* and *ServerSkeleton*, explained below), as well as the units' interconnections. A distribution architecture also determines how the particular requirements are addressed by individual units. The connector implementation at lower levels of the element hierarchy is defined recursively in such a way that each level of nesting is described by an *element architecture* specifying the horizontal composition of a specific element from its sub-elements (Fig. 6). Here, a sub-element is referred to just by its outer boundary – *element type*, which is to be further refined by an element architecture. This facilitates automated synthesis and reuse of hierarchically composable elements in a way similar to hierarchical components [BP04, CL02, TMD10]. For example, the unit *For\_Callee* is reflected as the *SerializedServerSkeleton* element (of type *ServerSkeleton*), the architecture refining its type introduces two sub-elements *SocketFactorySkeleton* and *CallSerializer*. In a similar vein, the architecture refining the *SocketFactorySkeleton*'s type defines two sub-elements *SocketFactoryProvider* and *RMISkeleton*.

Elements (including the top-level ones) interact via *ports*. A port can be either local or remote. A local port (e.g., in *FileLogger*) serves for internal communication of elements not directly participating in the distributed communication (based on local procedure calls). It is specified either as provided or required to emphasize where the communication is initiated in order allow more precise design of elements, facilitating the automated synthesis. A remote port (e.g., in *RMISub*) serves for distributed communication among units. In this case, the communication depends on

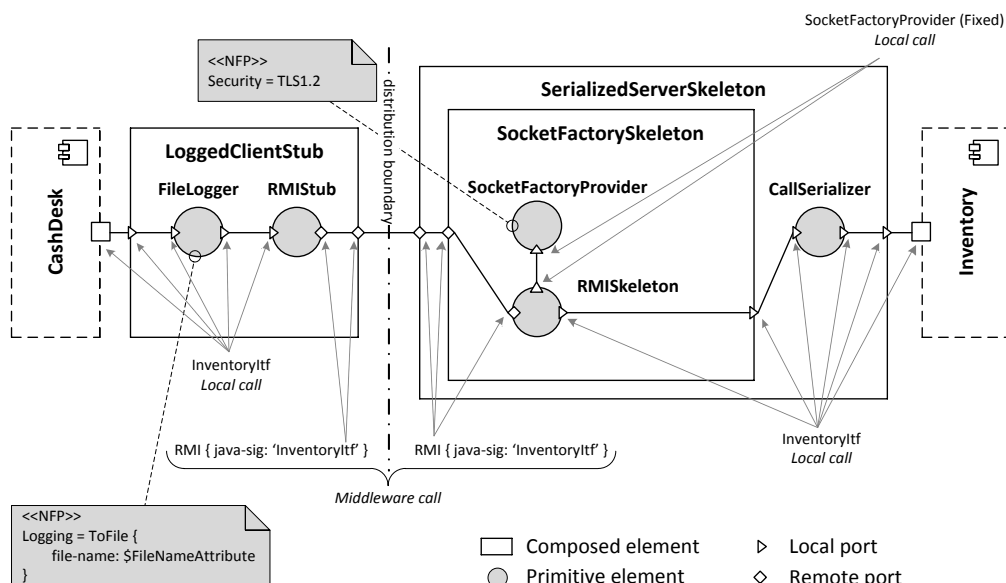


Fig. 6 Example of a full composition of a connector

the employed middleware. From the connector-designer perspective, the specifics of the middleware-based communication are intentionally abstracted. Port bindings are captured at the level of the parent element architecture and include both binding at the same level of nesting (e.g., `FileLogger` and `RMISub`) and port delegation (e.g., `SocketFactorySkeleton` and `RMISSkeleton`).

Eventually, every port is associated with a *signature*, which determines the interface type of the port. The association is either explicitly defined (e.g., the port signature of `SocketFactoryProvider`), or propagated – inferred either from the association of the other element’s ports bound to the current port (implicit propagation), or from the association of the other ports of the same element (signature constraints). In a particular connector implementation, the propagated signatures are typically associated with the signature of a component interface (e.g., `InventoryItf`). Using signature constraints, signature association can be propagated either directly (e.g., `FileLogger`), or as parameters of structured signatures (e.g., the remote interface of `RMISub`). In case of parameterized signatures, the implicit propagation implies that the parameterization of one signature has to match the parameterization of the other (e.g., `java-sig: ‘inventoryItf’`). This way, element architectures can be independent of the actual signatures (can use the signature propagation only) and are thus generic, which also facilitates reuse.

As for features, by following the idea of element composition, we assume every feature to be addressed by an element can be addressed either (i) directly by the element, or (ii) by one of its sub-elements. This allows isolation of features and facilitates the element composition. An example of (i) is the logging feature in the unit `For_Caller`. In this case, the entire task of logging is addressed by `FileLogger`, which intercepts and logs the communication going through the unit. An example of (ii) is the security feature delegated by `SocketFactorySkeleton` to `SocketFactoryProvider`. In general, a part of the feature value may be left unspecified in an element, since it is to be later on determined from requirements specification. This is captured by the element *attribute* concept, i.e., a parameter of the element (e.g., the file name attribute for logging to a file in `FileLogger`).

From the connector designer perspective, the *runtime-environment requirements* of an element architecture are expressed as constraints on dock capabilities. In other words, runtime-environment requirements are abstraction of a particular deployment.

In general, all the concepts of the design view are intended for reuse in multiple applications/connectors. For the purpose of this text, the design view of a connector is assumed to be described by an *artifact specification*.

## 2.1 Summary of the concepts and problem refinement

For the purpose of this paper, we will describe a particular connector instance using the sets and relations in Fig. 7, which shows key parts of the connector instance meta-model. In the rest of this paper, the description of a particular connector instance will be referred to as a connector instance configuration (CIC).

Specifically, all the key CIC concepts, i.e., connector, role, role instance, unit, element, sub-element, port, signature, feature, connector feature, attribute, and deployment dock, are represented explicitly via sets. However, the concepts of element type, communication style, and dock capabilities are not reflected explicitly, since they are rather a part of a design specification, being in a CIC “blended in” already. Further, the meta-model captures the key relations which are specific to a particular connector instance and its elements.

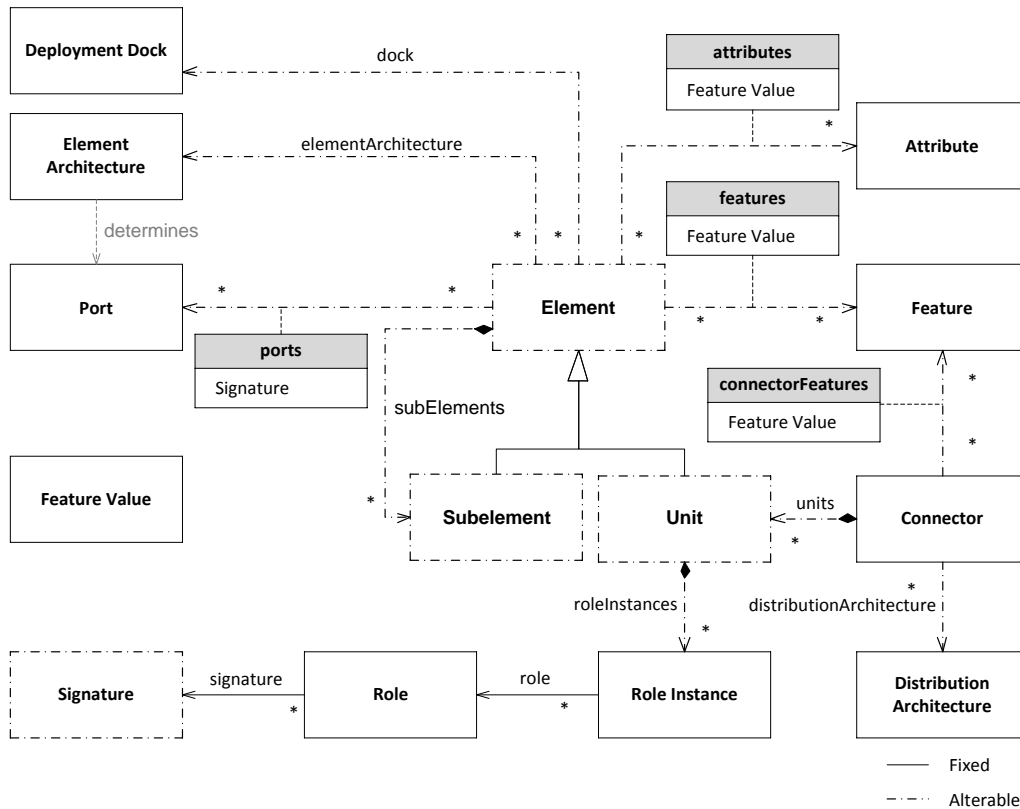


Fig. 7 Sets and relations describing a connector – meta-model

A particular connector (its CIC) is therefore an instance of the meta-model. The instance is determined by (a) an artifact specification (predefined), and by (b) requirements and deployment specifications (application-specific).

However, given both (a) and (b), while some of the sets and relations in the meta-model are fully determined (solid line), some are underspecified (dashed line), and therefore determined only partially (*alterable sets and relations*). In other words, multiple realizations of the alterable sets and relations can satisfy the given (a) and (b); i.e., multiple variants of the connector exist. For example, the variants can be introduced by the existence of several element architectures suitable for refining a sub-element. The alterable sets and relations are determined by the constraints inferred from (a) and (b). For the sake of brevity, we do not explicitly represent all of these constraints in the meta-model (the relation *determines* serves as an example). Even though neither the concept of element architecture nor distribution architecture is required in CIC, both of them are included in the meta-model in order to make it possible to express these constraints explicitly.

Overall, the problem of finding a CIC can be interpreted as finding a realization of the alterable sets and relations with respect to (a) and (b). If multiple alternatives exist, an optimal one should be chosen.

Specifically, a realization of the alterable sets and relations in CIC embodies the following:

- (i) a particular distribution architecture that conforms to the communication style specified in the requirements specification and definition of connector units;
- (ii) vertical composition of element architectures (each of them describing a horizontal composition of its sub-elements) determining, which element architecture to choose for each (sub-) element in a distribution architecture (recursively at all levels of element nesting);
- (iii) actual parameters for the element architectures by providing actual signatures for the elements' ports.

### 3. OVERVIEW OF THE ARCAS METHOD

Following the ideas of Section 2.1, automated resolution of CIC can be viewed as relational constraint solving problem where realization of the fixed sets and relations, as well as the constraints over the alterable sets and relations, form the constraint specification and a realization of the alterable sets and relations represents a solution to the problem. For this purpose, it is advantageous to employ a constraint-solving technique based on a modeling language rich enough to express the required concepts. In general, relational constraint solving languages such as relational logic are well suited to this purpose.

Informally put, we create a constraint specification, referred to as *connector theory* (CT) – in the sense of logic, capturing specification of a particular connector in terms of a logic theory, so that a model of such theory represents a CIC. More precisely, since a connector theory may have more than one model, this theory basically represents a description of a set of all the alternative CICs of the connector. In consequence, a model of a connector theory (in the sense of logic), provides representation of all the sets and relations of the CIC meta-model from Fig. 7, so that Fig. 7 can be also viewed as the meta-model of connector theory.

Expressed in the terms of the specifications, such a connector theory has the following important semantic properties:

(i) *Obeys the requirements specification.* This ensures that the connector captured by each model of the theory complies with a given requirements specification in the terms of communication style, signatures of its roles, and feature requirements.

(ii) *Obeys deployment specification.* This ensures that the connector units captured by each model of the theory are able to run in the target runtime platform nodes.

(iii) *Guarantees composition and refinement consistency.* This ensures that the elements in the connector architecture captured by a model of the theory are able to cooperate. This is important because not necessarily all assignments of sub-elements yield a working connector. A negative example would be to require combining RMI server skeleton with CORBA client stub.

Advantageously, CT can be constructed in an automated way by transformation of the specifications (Fig. 8). Specifically, CT consists of four parts based on: (i) a definition of the abstractions global to all CT theories (reflecting the meta-model), (ii) an image of the suitable element architectures with respect to the given dock capabilities (i.e., only the element architectures able to run in the deployment docks specified by deployment specification are considered), (iii) an image of the suitable distribution architectures with respect to the specified communication style, (iv) an image of the requirements imposed by connector requirements and deployment specifications (i.e., imposed on port signatures, features, deployment, etc.). Except for part (i), shared for all applications and created by hand in advance, all the other parts are produced by the transformations. As an aside, the representation of communication styles, element types, and deployment is subject to “inlining”. For example, in (ii) and (iii) each element type is inlined by a list of all the element architectures suitable for this type.

Once CT is produced, a constraint solver available for the selected modeling language is employed to find out a model of CT. This model is easily programmatically converted to the corresponding CIC.

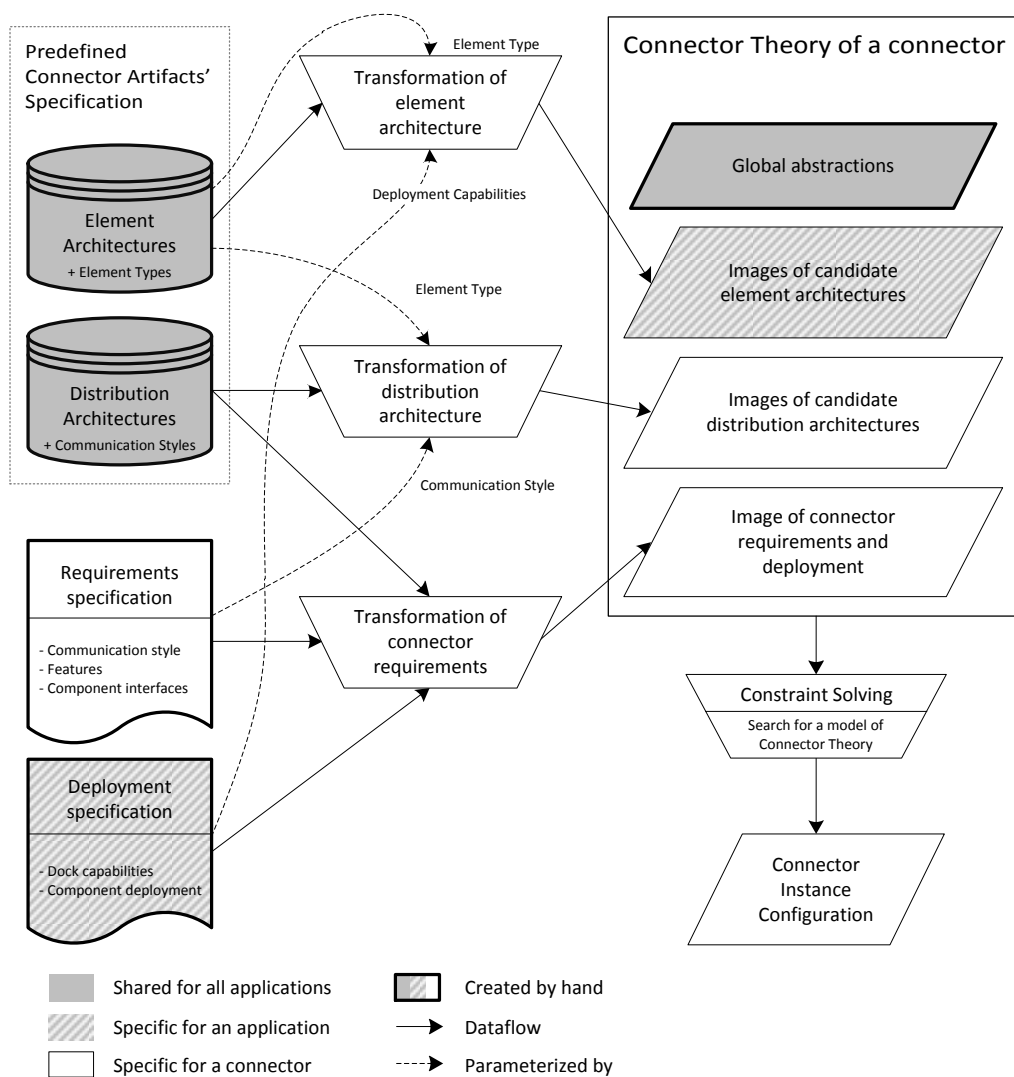


Fig. 8 Overview of the ARCAS method

The specifications are described in more detail in Section 4. The transformations of the specifications producing the connector theory are elaborated in Section 5. Since the Alloy modeling language [J02, J06] and its solver Alloy Analyzer are good candidates for representing CT (it provides convenient syntax for definition of relations and their constraints), in Section 6 we describe the representation of CT using Alloy.

Referring back to Section 1, it should be emphasized that ARCAS is intended to be applied at either the deployment or runtime stage of an application. In the latter case, this would be due to a runtime modification of the architecture and/or deployment of the application.

#### 4. ARCAS INPUT: SPECIFICATIONS

In this section, we will elaborate on the specifications required as the input of the transformations in ARCAS (Fig. 8); the specifications were conceptually outlined in Section 2. In principle, each of them is defined by its meta-model, i.e., abstract syntax, and for practical reasons ARCAS includes connector definition language (CDL), i.e., concrete syntax, in which we will provide examples.

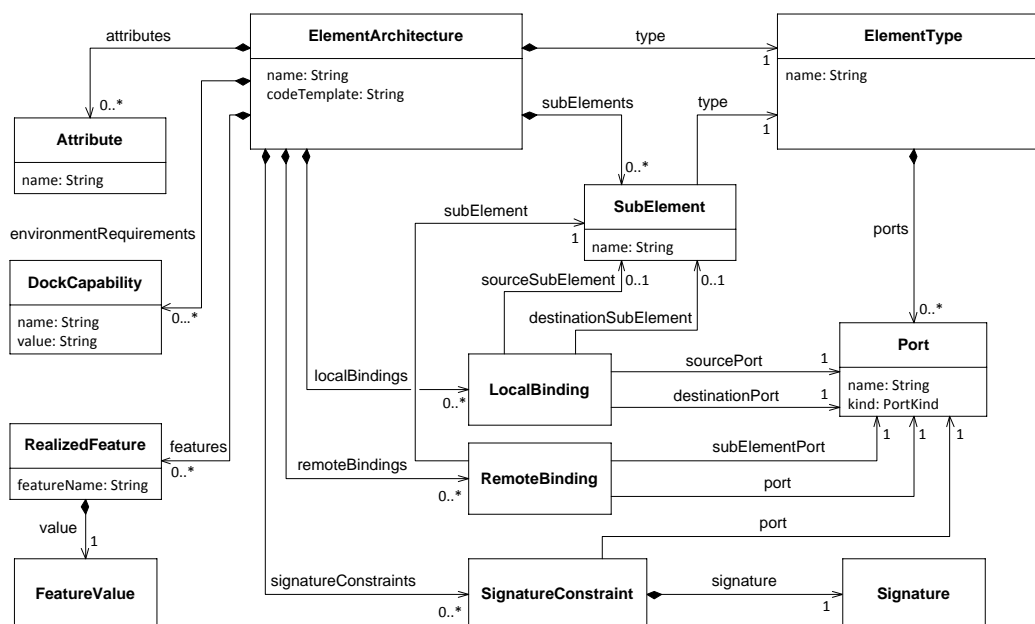


Fig. 9 Abstract syntax of Element Architecture specification

We will fully describe the abstract syntax and semantics, and give examples of the element type and element architecture specifications. For brevity, the other specifications (i.e., communication style, distribution architecture, connector requirements, and deployment specifications) are illustrated by an example based on Fig. 4 and Fig. 6, while their abstract syntax is provided in Appendix A.

#### 4.1 Element Type & Element Architecture Specifications

The basic abstraction of `ElementArchitecture` (Fig. 9) is `ElementType`, which defines the external interface of an element, i.e., its ports. A port is defined as provided, required, or remote. The following illustrates specification of `ClientStub` and `Logger` element types (the types of `LoggedClientStub` and `FileLogger` from Fig. 6) in CDL:

```

element-type ClientStub
  local-provides: callIn      -- <port kind>: <port name>
  remote: mw

```

```

element-type Logger
  local-provides: callIn
  local-requires: callOut

```

An element architecture refines an element type by determining a hierarchical composition of elements at two adjacent levels of nesting (the lower level is determined by the `subElements` relation in Fig. 9). Thus, only horizontal composition (of the elements at the lower level) is specified for `ElementArchitecture`; as to vertical composition, it is here expressed by enforcing a particular element type for a sub-element.

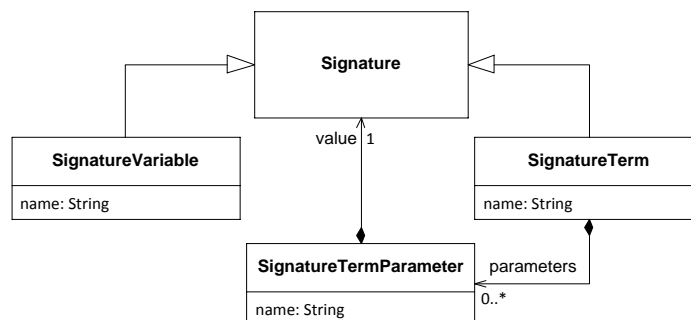


Fig. 11 Abstract syntax for Signature specification

Further, indication of port bindings is defined separately for pairs of provided and

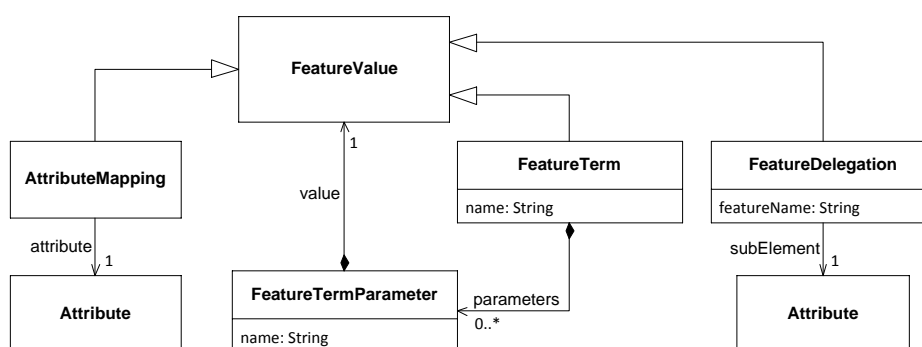


Fig. 10 Abstract syntax for specification of features

required ports (*LocalBinding*), and for delegation among remote ports (*RemoteBinding*). The former case expresses either delegation or communication at the same level of nesting.

Port signature propagation is captured by expressing relations between the ports of a single element architecture (*SignatureConstraint*); these relations enforce equal or equally parameterized signatures of the related ports. This ensures an abstraction over element implementation, providing only the information important for vertical composition. The way *Signature* is defined in Fig. 11 indicates that a signature takes the form of either a signature variable, acting as a signature parameter, or a signature term (possibly parameterized).

Finally, as for the features (Fig. 10) realized by the element architecture (*RealizedFeature*), a feature value is represented by a possibly structured feature term. The value of a feature (sub-) term is either delegated to a sub-element (*FeatureDelegation*), explicit (set explicitly via a fixed feature term), or parameterized by an *attribute* (*AttributeMapping*), playing a similar role as signature variables in signature terms. Assuming an element  $e$ , the actual value of its attribute is defined either in a parent element (recursively) or directly at the level of the whole connector (in the requirements specification); in both cases the value is propagated to  $e$  by feature delegation. The feature delegation requires that higher-level element architectures need to anticipate the features of their sub-elements, which is in principle an extensibility issue; potential solutions are discussed in Section 6.

Element architecture specification also includes the runtime-environment requirements as introduced in the OMG D&C specification [OMG04, BB04]. This reflects the necessary runtime support for deployment of the element architecture. For

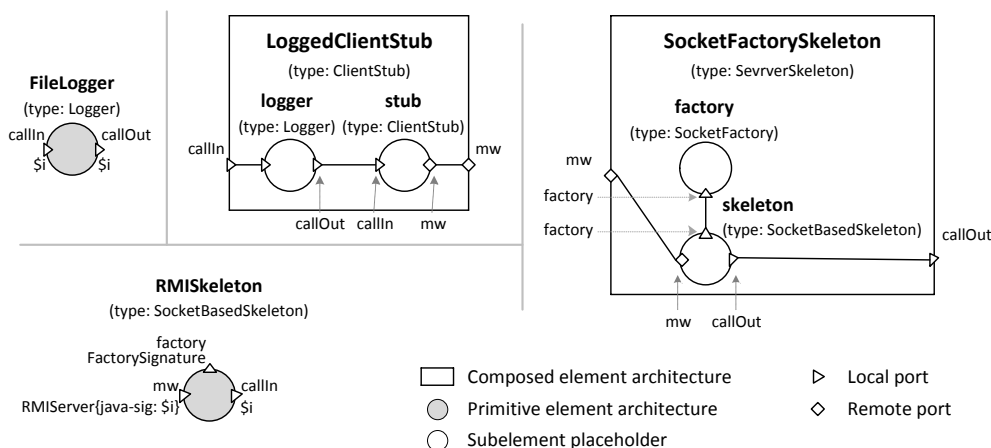


Fig. 12 Elaboration of FileLogger, RMISkeleton, LoggedClientStub, and SocketFactorySkeleton element architectures from Fig. 6

simplicity, we represent these requirements by means of name-value pairs (DockCapability).

The CDL specification of the LoggedClientStub, FileLogger, and SocketFactorySkeleton element architectures from Fig. 12 (elaboration from Fig. 6) can take the following form:

```

element-architecture LoggedClientStub
  of-type: ClientStub
  sub-elements:
    stub: ClientStub           -- <name>: <type>
    logger: Logger
  bindings:
    this.callIn -> logger.callIn  -- -> local binding
    logger.callOut -> stub.callIn  -- <subelement>.<port>
    this.mw <-> stub.mw          -- <-> remote binding
  features:
    logging -> logger.logging     -- feature delegation

element-architecture FileLogger
  of-type: Logger
  signature-propagation:
    callIn: $I                  -- $I is a signature variable
    callOut: $I
  features:
    logging: ToFile{            -- explicit feature, set to ToFile{...}
      name: $FileName           -- 'name' is a feature term parameter
    }                            -- $FileName is an attribute
  environment-requirements: {"Java VM" -> "JDK 1.4" }

element-architecture SocketFactorySkeleton
  of-type: ServerSkeleton
  sub-elements:
    skeleton: SocketBasedSkeleton  -- <name>: <type>
    socket-factory: SocketFactory
  bindings:
    skeleton.factory -> socket-factory.factory  -- -> local binding
    skeleton.callOut -> this.callOut           -- <sub-element>.<port>
    this.mw <-> skeleton.mw                   -- <-> remote binding
  features:
    security -> socket-factory.security        -- feature delegation

element-architecture RMISkeleton
  of-type: SocketBasedSkeleton
  signature-propagation:

```

```

callOut: $I          -- $I is a signature variable
mw: RMI{ java-sig: $I } -- parameterized signature
                    -- 'java-sig' is a signature parameter
factory: FactorySignature -- fixed port signature
environment-requirements: {"Java VM" -> "JDK 1.6" }

```

Here, an important part is the definition of features and attributes, element types, signature propagation, and environment requirements:

FileLogger and logging give an example of an explicit feature. The feature value is parameterized by the FileName attribute, which is used in the implementation of FileLogger. The security feature of SocketFactorySkeleton and logging of LoggedClientStub illustrate the delegation of features. Here, the meaning is that SocketFactorySkeleton provides security only if this feature is provided by the socket-factory sub-element (in the positive case, socket-factory also defines the value of security).

As to element types, notice that both LoggedClientStub and its sub-element stub are of the same element type. Thus, both LoggedClientStub and RMISStub can be used to refine the For\_Caller unit. In this example, LoggedClientStub was preferred over RMISStub because of its logging feature.

Signature propagation is illustrated by FileLogger, where the use of the same signature variable \$I for both callIn and callOut indicates that both ports have to have the same signature. A more complex example is RMISkeleton, where the signature of the callOut port is propagated as the java-sig parameter of the mw port's signature. The RMISkeleton specification also illustrates definition of a fixed port signature FactorySignature.

The environment requirements are expressed by means of required dock capability values. RMISkeleton requires the target deployment dock to support JDK 1.6.

## 4.2 Remaining Specifications

4.2.1 Communication Style. As outlined in Section 2, specification of a communication style determines the connector roles and their cardinality. The communication style of the connector from Fig. 3 can be specified as:

```

communication-style ProcedureCall
roles:
  Callee          -- lower bound = 1, upper bound = 1
  Caller(0..n)   -- lower bound = 0, upper bound = n

```

4.2.2 Connector Requirements. The key idea of connector requirements is that the particular connector architecture to be used for the connector is not specified explicitly; instead, it is declaratively determined by the selected communication style, required features, and desired deployment of components.

As requirements specification is application-specific, the components to be connected and their interfaces (including signatures) have to be defined. Based on this, connector requirements are specified by selecting a communication style, mapping roles to the component interfaces (i.e., defining connector endpoints), and by defining the required features of the connector. Feature requirements are determined by enumerating the acceptable feature values, where a value is represented by a feature term. There is also the option to define a feature requirement for a particular endpoint. Feature requirements can be composed using propositional operators.

The specification of connector requirements for the example from Fig. 3 can take the following form:

```

component CashDesk
  requires: InventoryItf inventory -- <interface kind>: <signature> <name>

component Inventory
  provides: InventoryItf inventory

```

```

connector CashDesk_to_Inventory
  communication-style: ProcedureCall
  endpoints:
    CashDesk.inventory as Caller -- <component.interface> as <role>
    Inventory.inventory as Callee
  features:
    security in { TLS1.2, SSL3 } -- <feature> in {<possible values>}
    Inventory.inventory.logging in { ToFile { name: "inventory.log" }}
                                   -- <component>.<interface>.<feature>
                                   -- 'name' is feature term parameter

```

4.2.3 Deployment. A deployment specification expresses the desired deployment of the connected components (not of connector deployment!). It determines the allocation of components to deployment docks and the capabilities of the deployment docks. The dock capabilities are based on the OMG D&C [OMG04, BB04] standard (syntactically expressed by means of typed name-value pairs). The specification of deployment for the example from Fig. 4 can take the following form:

```

allocation
  CashDesk to DockA
  Inventory to DockB

dock DockA
  capabilities: {"Java VM" -> "JDK 1.4", "OS" -> "Linux 2.4.28"}

dock DockB
  capabilities: {"Java VM" -> "JDK 1.6", "OS" -> "Linux 2.6.36"}

```

4.2.4 Distribution Architecture. The distribution architecture specification determines the desired communication style, the element type of the top-level element architectures refining units, their cardinality, the units' remote bindings, and mapping of roles to ports in units. The specification also explicitly states how features are addressed by specific units (feature delegation). The distribution architecture specification of the connector from Fig. 5 can take the following form:

```

distribution-architecture RPC
  communication-style: ProcedureCall
  units:
    For_Callee: ServerSkeleton -- <name>: <element type>
    For_Caller (0..n): ClientStub -- <name>(<cardinality>): <type>
  remote-bindings:
    For_Caller(i).mw <-> For_Callee.mw -- <unit>.<port>
  role-mapping:
    For_Caller(i).callIn as Caller(i) -- <unit>.<port> as <role>
    For_Callee.callOut as Callee
  features:
    security -> For_Callee.security -- feature delegation

```

An important concept is the specification of multiplicity of unit instances such as `For_Caller (0..n): ClientStub`, which says that there will be up to  $n$  `For_Caller` units of the type `ClientStub`. In consequence, multiplicity is to be expressed also in `remote-bindings` and `role-mapping`. For example, `For_Caller(i).mw <-> For_Callee.mw` means that the `mw` port of all the `For_Caller` units is bound to the `mw` port of the `For_Callee` unit. Likewise, `For_Caller(i).callIn as Caller(i)` means that the `callIn` port of each `For_Caller` unit has the role `Caller`.

## 5. CONSTRUCTING CONNECTOR THEORY BY TRANSFORMATIONS

As outlined in Section 3, a connector theory can be constructed in an automated way by transformations of the specifications. In this section, we describe the transformations and their products in terms of propositional logic with relational

calculus, using only basic logical and relational operators. In general, a transformation results in a formula. The set of all formulas resulting from transformations of the specifications relevant to a particular connector forms the corresponding connector theory.

The formulas express the constraints on alterable sets/relations by binding them to the fixed ones (Fig. 7). In the formulas, we rely on predicates, the semantics of which is explained informally in this section to help the reader get an intuitive understanding of them; detailed semantics is provided only for a selection of the predicates. We further elaborate on the description of selected predicates in Section 6 by implementing them in Alloy<sup>1</sup>. There, we also illustrate a construction of the fixed sets and relations. For the fixed sets and relations, we assume that a realization is available, implied by the specifications (both those of predefined artifacts, and those of requirements and deployment, specific to a connector).

For defining the actual transformations based on predicate logic with relational algebra, it is advantageous to introduce a slightly modified version (Fig. 13) of the connector meta-model from Fig. 7. The rationale for such modification is that certain relations of the original meta-model (i.e., *features*, *connectorFeatures*, *attributes*, and *ports*) are parameterized. In order to capture this in relational calculus, the originally binary parameterized relation is now expressed as a ternary relation (e.g., the relation *features* in Fig. 13). In addition, given an element *e*, in order to express constraints over its particular sub-element *se*, the name of the sub-element is explicitly employed. This is necessary, since the constraints on *se* are imposed in the element architecture of *e*. Therefore, the original binary relation *subElements* is replaced by a

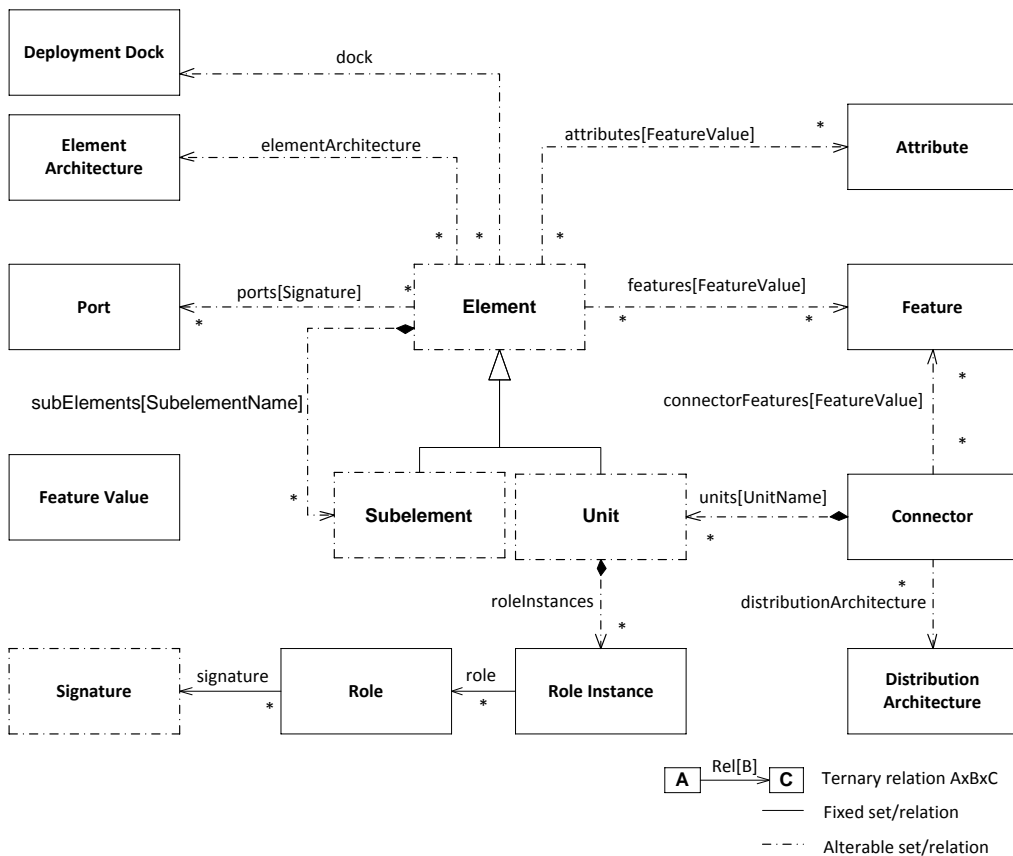


Fig. 13 Modified connector theory meta-model (by ternary relations)

<sup>1</sup> The full formalization in Alloy is provided at [http://d3s.mff.cuni.cz/projects/components\\_and\\_services/arcas/](http://d3s.mff.cuni.cz/projects/components_and_services/arcas/)

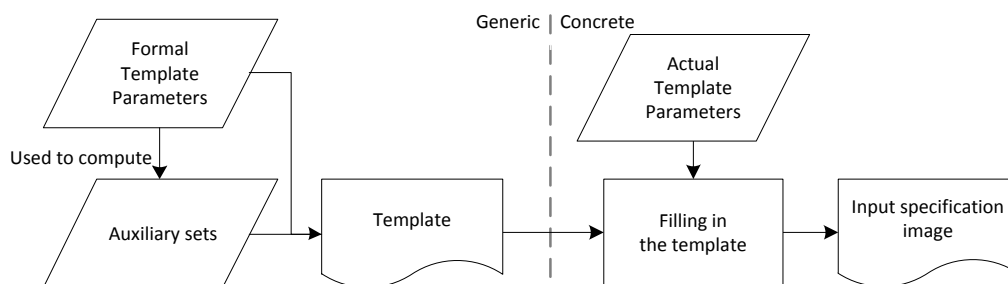


Fig. 14 Generic workflow of transformations

ternary relation introducing an identifier of the sub-element (the identifier is taken over from the definition of  $se$  in the element architecture of  $e$ ). The relation  $units$  is to be modified in a similar way to  $subElements$ .

Because the formulas resulting from a single transformation have a fixed internal structure, we present the transformations in terms of parameterized templates (Fig. 14). The template parameters are derived from the transformation parameters (Section 3). Based on actual transformation parameters, the transformation produces a specification's image (i.e., a part of the connector theory under construction) by filling in the template.

In detail, we describe the transformation of a composite element architecture specification focusing on the corresponding template and derivation of the template parameters. For brevity, the effect of other transformations (i.e., of primitive element architecture, distribution architecture, and requirements specifications) is illustrated by an example of their product, while their full description is provided in Appendix B.

## 5.1 Transforming specification of a primitive Element Architecture

The transformation workflow for a primitive element architecture specification is illustrated in Fig. 15.

5.1.1 Template. The template (Fig. 16) formulates the constraints on any element  $e$  employing a particular primitive element architecture  $ea$ , the specification of which is subject to the transformation. The template takes the form of an implication, stating that  $e$  employing  $ea$  has to conform to a conjunction of constraints; these are derived from the semantics of the primitive element architecture concept (Section 2 and 4.1).

Technically, in the meta-model these constraints restrict the relations which are (possibly transitively) related to  $e$  as an element of the set  $Element$ . The antecedent of the implication assumes validity of the  $ElementHasEA$  predicate over the  $elementArchitecture$  relation ( $e$  employs  $ea$ ). In the consequent, it is necessary to express that deployment of  $e$  is possible only to the docks compatible with  $ea$  ( $AllowedDocks$  set); this is reflected by constraining the relation  $node$  by  $AllowedDeployment$ . Further,  $e$  can feature only the ports defined by the element type of  $ea$  (the  $ports$  relation is constrained by  $AvailablePorts$ ). It is also desirable to make

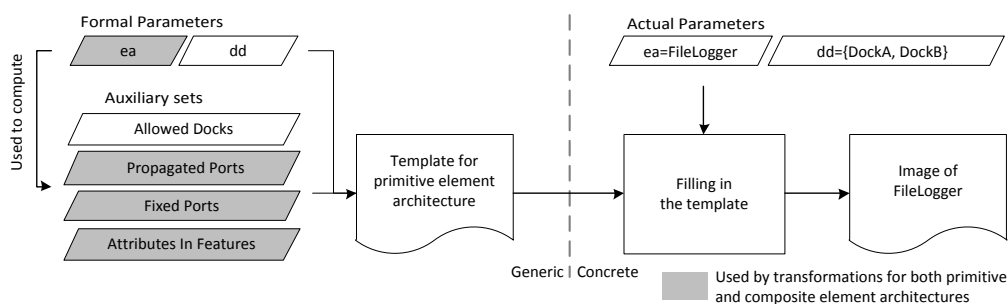


Fig. 15 Transformation workflow for primitive element architecture specifications

```

 $\forall e \in Elements : ElementHasEA(e, \langle ea \rangle) \Rightarrow ($ 
  AllowedDeployment(e,  $\langle AllowedDocks \rangle$ )  $\wedge$ 
  AvailablePorts(e,  $\langle ea.type.ports \rangle$ )  $\wedge$ 
  SubElements(e,  $\emptyset$ )  $\wedge$ 
  SupportedFeatures(e,  $\langle ea.features \rangle$ )  $\wedge$ 
  RequiredAttributes(e,  $\langle ea.attributes \rangle$ )  $\wedge$ 

  -- signature propagation
  <foreach (port1, port2) in PropagatedPorts >
    PropagateSignature(e,  $\langle port1 \rangle$ ,  $\langle port2 \rangle$ )  $\wedge$ 
  <endforeach>

  -- enforcing fixed signatures
  <foreach (port1, port2) in FixedPorts >
    PortHasSignature(e,  $\langle port \rangle$ ,  $\langle signature \rangle$ )  $\wedge$ 
  <endforeach>

  -- enforcing values of the realized features
  <foreach f in ea.features >
    FeatureHasValue(e,  $\langle f.name \rangle$ ,  $\langle Encode(f.value) \rangle$ )  $\wedge$ 
  <endforeach>

  -- defining attributes referred in feature values
  <foreach (feature, parameter, attribute) in AttributesInFeatures>
    FeatureValueUsesAttribute(e,  $\langle feature \rangle$ ,  $\langle parameter \rangle$ ,  $\langle attribute \rangle$ )  $\wedge$ 
  <endforeach>
)

```

Fig. 16 Template for image of primitive element architecture

sure that the set of sub-elements of  $e$  is empty since  $ea$  is primitive (via constraining the *subElements* relation by *SubElements*). As for the further constraints, only the realized features and declared attributes by  $ea$  may be employed by  $e$  (*features* and *attributes* are constrained by *SupportedFeatures* resp. *RequiredAttributes*).

According to the signature propagation constraints for  $ea$ , the signatures/signature parameters of the involved ports of  $e$  have to be equal (*ports* is constrained by *PropagateSignature*). Note that we have abstracted away details of signature propagation for the sake of brevity. In addition, it is necessary to enforce the fixed port signatures of  $ea$  (*ports* is constrained by *PortHasSignature*). Similar to the fixed signatures, the feature values defined in  $ea$  have to be propagated to  $e$  (via constraining *features* by *FeatureHasValue*). Since a feature value can use an attribute to valuate its parameter, the attribute value and the respective feature parameter value have to be equal (*features* and *attributes* are constrained by *FeatureValueUsesAttribute*). The semantics of selected predicates is presented in Table 1 via both textual description and logical formula.

Table 1. Semantics of predicates for primitive element architecture image

| Predicate name                      | Description                                                                                                                                                       | Formula                                                                                                |
|-------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------|
| ElementHasEA<br>( $e, ea$ )         | $e$ is in the <i>elementArchitecture</i> relation with $ea$                                                                                                       | $(e, ea) \in element - Architecture$                                                                   |
| AllowedDeployment<br>( $e, nodes$ ) | $e$ can be in the <i>node</i> relation only with the elements of <i>nodes</i>                                                                                     | $\forall n \in DeploymentNode:$<br>$(e, n) \in node \Rightarrow n \in nodes$                           |
| AvailablePorts<br>( $e, portSet$ )  | $e$ can be in the <i>ports</i> relation only with the elements of <i>portSet</i> , no matter the signature (it is a ternary relation of element, port, signature) | $\forall p \in Port, \forall s \in Signature:$<br>$(e, s, p) \in ports$<br>$\Rightarrow p \in portSet$ |

|                                                                                            |                                                                                                                                                                                                                                                                                                                                                                                                   |                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RequiredAttributes<br>( $e$ , $attrSet$ )                                                  | $e$ is in the <i>attributes</i> relation with exactly the elements of <i>attrSet</i> , no matter the attribute value (it is a ternary relation of element, attribute, attribute value)                                                                                                                                                                                                            | $\forall a \in Attribute,$<br>$\forall v \in AttributeValue:$<br>$(e, v, a) \in attributes$<br>$\Leftrightarrow a \in attrSet$                                                                                          |
| FeatureValueUses-<br>Attribute<br>( $e$ ,<br>$feature$ ,<br>$parameter$ ,<br>$attribute$ ) | the feature value $fv$ , which is in the <i>features</i> relation with $e$ and <i>feature</i> , and the attribute value $av$ , which is in the <i>attributes</i> relation with $e$ and <i>attribute</i> , are in the <i>parameter</i> relation (representing the particular feature parameter) Here, a feature parameter is modeled using a dedicated relation (not reflected in the meta-model). | $\forall fv \in FeatureValue,$<br>$\forall av \in FeatureValue:$<br>$\{(e, fv, feature) \in features$<br>$\wedge$<br>$(e, av, attribute) \in attributes\}$<br>$\Rightarrow \Leftrightarrow$<br>$(fv, av) \in parameter$ |

5.1.2 Formal parameters. The formal parameters of the primitive-element-architecture transformation template are: (i) the element architecture  $ea$  to be transformed and (ii) the actual deployment specification  $dd$  (Fig. 15).

5.1.3 Auxiliary sets. These are *AllowedDocks*, *PropagatedPorts*, *FixedPorts*, and *AttributesInFeatures* (Fig. 15). *AllowedDocks* contains all the deployment docks in which  $ea$  is allowed to run. *PropagatedPorts* comprises all the pairs of ports of  $ea$  that use the same signature parameter and therefore are involved in signature propagation (note that we have again abstracted away the details of signature propagation for the sake of brevity). *FixedPorts* includes all the (*port*, *signature*) pairs such that *signature* of *port* is explicitly defined using a fixed signature term (i.e., a signature term without delegation or signature variables). *AttributesInFeatures* contains all the tuples (*feature*, *feature parameter*, *attribute*) such that *attribute* is used as the value of *feature parameter* which is part of the feature term associated with *feature*. Mathematically, the sets can be defined by the expressions in Fig. 17.

5.1.4 Example: Image of FileLogger. Referring back to the example from Sections 2 and 4, consider the actual parameters  $ea = FileLogger$ ,  $dd = \{DockA, DockB\}$ , and  $eas = \{FileLogger, \dots\}$ ; the transformation will produce the image presented in Fig. 18.

## 5.2 Transforming specification of a composite Element Architecture

The transformation workflow for a composite element architecture specification is illustrated in Fig. 19.

5.2.1 Template. The template (Fig. 20) is very similar to the template for primitive

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| $AllowedDocks = \{dock.name   dock \in dd.docks \wedge$ $OMGD\&C\_Compatible(d.capabilities, ea.envRequirements)\}$ $PropagatedPorts = \{(port1, port2)   port1, port2 \in ea.type.ports \wedge$ $sc1, sc2 \in ea.signatureConstraints \wedge$ $sc1.port = port1 \wedge sc2.port = port2 \wedge$ $UseTheSamePlaceholder(sc1.signature, sc2.signature)\}$ $FixedPorts = \{(p, sc.signature)   p \in ea.type.ports \wedge$ $sc \in ea.signatureConstraints \wedge$ $sc.port = p \wedge IsFixedTerm(sc.signature)\}$ $FixedFeatures = \{f   f \in ea.features \wedge IsFixedFeatureTerm(f.value)\}$ $AttributesInFeatures = \{(f.name, param.name, a.name)   a \in ea.attributes \wedge$ $f \in ea.features \wedge$ $term \in TransitiveClosure(f.value) \wedge$ $param \in term.parameters \wedge$ $param.value = a\}$ |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Fig. 17 Auxiliary sets for primitive element architecture

```

 $\forall e \in \text{Elements} : \text{ElementHasEA}(e, \text{FileLogger}) \Rightarrow ($ 
   $\text{AllowedDeployment}(e, \{\text{DockA}\}) \wedge$ 
   $\text{AvailablePorts}(e, \{\text{CallIn}, \text{CallOut}\}) \wedge$ 
   $\text{SubElementEAs}(e, \emptyset) \wedge$ 
   $\text{SupportedFeatures}(e, \{\text{Logging}\}) \wedge$ 
   $\text{RequiredAttributes}(e, \{\text{FileName}\}) \wedge$ 

  -- signature propagation
   $\text{PropagateSignature}(e, \text{CallIn}, \text{CallOut}) \wedge$ 

  -- enforcing fixed signatures

  -- enforcing values of the realized features
   $\text{FeatureHasValue}(e, \text{Logging}, \text{LoggingToFile}) \wedge$ 

  -- defining attributes referred in feature values
   $\text{FeatureValueUsesAttribute}(e, \text{Logging}, \text{fileName}, \text{FileName})$ 
 $)$ 

```

Fig. 18 Example: Image of FileLogger primitive element architecture

element architecture (i.e., it formulates constraints on element  $e$  employing  $ea$  in the form of an implication). Therefore, we will skip description of the common parts by referring the reader back to Section 5.1.

As for the constraints in the consequent of the implication (similar to the case of primitive element architecture), after enforcing that  $e$  can feature only the ports defined by the element type of  $ea$  ( $\text{AvailablePorts}$ ), the set of sub-elements of  $e$  is constrained so that it conforms to  $ea$  ( $\text{SubElements}$ ). In addition, it is desirable to make sure that only the realized features and declared attributes of  $ea$  may be employed by  $e$  ( $\text{SupportedFeatures}$  resp.  $\text{RequiredAttributes}$ ).

In contrast to primitive element architecture, a large part of the template focuses on sub-element properties and bindings. The definitions of the sub-elements are reflected as follows. For each sub-element, the set of element architectures that can be employed for its refinement is explicitly prescribed (constraining the  $\text{subElements}$  and  $\text{elementArchitecture}$  relations by  $\text{AllowedEAsForSubElement}$ ). This ensures the refinement consistency of the sub-elements. The deployment of the sub-elements is propagated to  $e$ , since  $e$  has to be deployed on the same node as its sub-elements (the relations  $\text{subElements}$  and  $\text{node}$  are constrained by  $\text{PropagateDeployment}$ ). Further, similar to  $e$  itself, it is desirable to reflect that a sub-element can feature only the ports defined by its element type prescribed in  $ea$  ( $\text{AvailablePorts}$ ). Then, the port bindings and port delegations are reflected by forcing the signatures of the involved ports to be equal (via constraining the  $\text{subElements}$  and  $\text{ports}$  relations by  $\text{PortBinding}$  resp.  $\text{PortDelegation}$ ). This, along with enforcing signature propagation, ensures the composition consistency of the sub-elements. Finally, after capturing fixed signatures

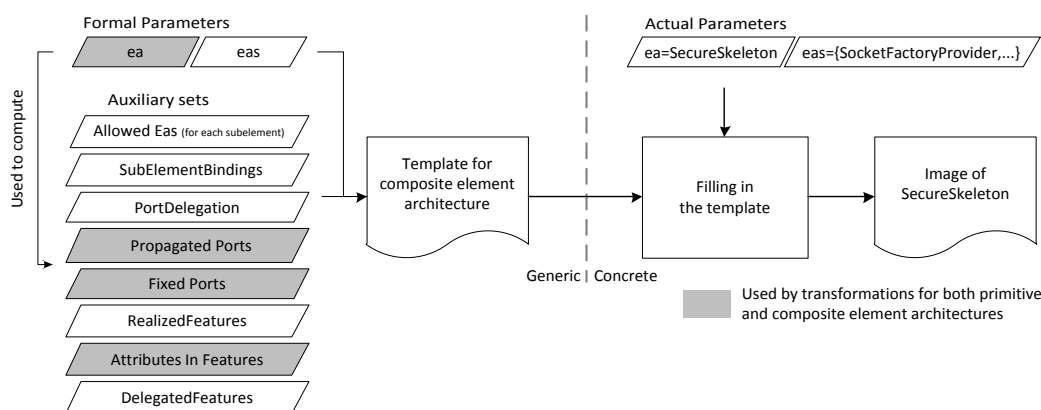


Fig. 19 Transformation workflow for composite element architecture

```

 $\forall e \in Elements : ElementHasEA(e, \langle ea \rangle) \Rightarrow ($ 
  AvailablePorts( $e, \langle ea.type.ports \rangle$ )  $\wedge$ 
  SubElements( $e, \langle ea.subElements \rangle$ )  $\wedge$ 
  SupportedFeatures( $e, \langle ea.features \rangle$ )  $\wedge$ 
  RequiredAttributes( $e, \langle ea.attributes \rangle$ )  $\wedge$ 

  -- definition of sub-elements
  <foreach se in ea.subElements >
    -- <se> sub-element
    AllowedEAsForSubElement( $e, \langle se \rangle, \langle AllowedEAs_{se} \rangle$ )  $\wedge$ 
    PropagateDeployment( $e, \langle se \rangle$ )  $\wedge$ 
    AvailablePorts( $e, \langle se \rangle, \langle se.type.ports \rangle$ )  $\wedge$ 
  <endforeach>

  -- port bindings
  <foreach (se1, port1, se2, port2) in SubElementBindings >
    PortBinding( $e, \langle se1 \rangle, \langle port1 \rangle, \langle se2 \rangle, \langle port2 \rangle$ )  $\wedge$ 
  <endforeach>
  /* port delegation */
  <foreach (port, se, seport) in PortDelegations>
    PortDelegation( $e, \langle port \rangle, \langle se \rangle, \langle seport \rangle$ )  $\wedge$ 
  <endforeach>

  -- signature propagation
  ...
  -- enforcing fixed signatures
  ...
  -- enforcing values of the directly realized features
  ...
  -- defining attributes referred in feature values
  ...
  -- feature delegation
  <foreach (feature, se, sefeature) in DelegatedFeatures >
    FeatureDelegation( $e, \langle feature \rangle, \langle se \rangle, \langle sefeature \rangle$ )  $\wedge$ 
  <endforeach>
)

```

Fig. 20 Template for image of composite element architecture

and signature propagation, as well as reflecting values of the directly realized features (the *RealizedFeatures* set) and attributes, in the same way as for primitive element architectures, the delegated features have to be captured (*subElements* and *features* are constrained by *FeatureDelegation*). The semantics of selected predicates is presented in Table 2 via both textual description and logical formula.

Table 2. Semantics of predicates for composite element architecture image

| Predicate name                                       | Description                                                                                                                                                                   | Formula                                                                                                                                                                                                 |
|------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AllowedEAsFor-SubElement<br>( $e, subelement, eas$ ) | the element $se$ , determined from <i>subElements</i> by $e$ and <i>subelement</i> , can be in the <i>elementArchitecture</i> relation only with the elements of <i>eas</i> . | $\forall se \in Element:$<br>$(e, subelement, se) \in subElements$<br>$\Rightarrow ($<br>$\forall ea \in ElementArchitecture:$<br>$(se, ea) \in elementArchitecture$<br>$\Rightarrow ea \in eas$<br>$)$ |
| PropagateDeployment<br>( $e, subelement$ )           | $e$ is in the relation <i>node</i> with the same deployment nodes as <i>subelement</i> .                                                                                      | $\forall se \in Element:$<br>$(e, subelement, se) \in subElements$<br>$\Rightarrow ($<br>$\forall n \in DeploymentNode:$<br>$(e, n) \in nodes \Leftrightarrow (se, n) \in nodes$<br>$)$                 |

|                                                                                 |                                                                                                                                                                                                                                                                                                         |                                                                                                                                                                                                                                                                            |
|---------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| PortBinding<br>( <i>e</i> ,<br>subelement1,<br>port1,<br>subelement2,<br>port2) | having the elements <i>se1</i> and <i>se2</i> , determined from <i>subElements</i> by <i>e</i> and <i>subelement1</i> and <i>subelement2</i> , respectively, <i>se1</i> and <i>port1</i> are in the relation <i>ports</i> with the same signatures as <i>se2</i> and <i>port2</i> .                     | $\forall se1, se2 \in Element,$<br>$(e, subelement1, se1) \in subElements$<br>$\wedge$<br>$(e, subelement2, se2) \in subElements$<br>$\Rightarrow$ (<br>$\forall s \in Signature:$<br>$(se1, s, port1) \in ports$<br>$\Leftrightarrow$<br>$(se2, s, port2) \in ports$<br>) |
| FeatureDelegation<br>( <i>e</i> ,<br>feature,<br>subelement,<br>se-feature)     | having the element <i>se1</i> , determined from <i>subElements</i> by <i>e</i> and <i>subelement</i> , a feature value <i>fv</i> is in the relation <i>features</i> with <i>e</i> and <i>feature</i> if and only if <i>fv</i> is in the <i>features</i> relation with <i>se</i> and <i>se-feature</i> . | $\forall se \in Element:$<br>$(e, subelement, se) \in subElements$<br>$\Rightarrow$ (<br>$\forall fv \in FeatureValue:$<br>$(e, fv, feature) \in features$<br>$\Leftrightarrow$<br>$(se, fv, se-feature) \in features$<br>)                                                |

5.2.2 Formal parameters. The formal parameters are: (i) the element architecture *ea* to be transformed and (ii) the set of all predefined element architectures *eas* (Fig. 19).

5.2.3 Auxiliary sets. These are *AllowedEAs<sub>se</sub>* (separately for each sub-element of *ea*), *SubElementBindings*, *PortDelegation*, *RealizedFeatures*, *AttributesInFeatures*, and *DelegatedFeatures* (Fig. 19). *AllowedEAs<sub>se</sub>* contains the element architectures that may be used for refining the sub-element *se* (decision is based on the element type of *se* defined by *ea*). *SubElementBindings* comprises all pairs of sub-elements of *ea* and their ports for which *ea* defines a binding. *PortDelegation* contains the tuples (*ea* port *port*, sub-element *se*, sub-element port *seport*) that participate in port delegation (including delegation of both local and remote ports). *RealizedFeatures* includes all the features realized directly by *ea*. *AttributesInFeatures* is exactly the same as in Section 5.1.3. *DelegatedFeatures* contains all the tuples (*ea* feature *f*, sub-

|                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| $AllowedEAs_{se} = \{subea \mid subea \in eas \wedge subea.type = se.type\}$<br>$SubElementBindings = \{(se1, port1, se2, port2) \mid se1, se2 \in ea.subElements \wedge$<br>$port1 \in se1.type.ports \wedge port2 \in se2.type.ports \wedge$<br>$\exists lb \in ea.localBindings \wedge$<br>$IsSource(lb, se1, port1) \wedge IsDestination(lb, se2, port2)\}$<br>$PortDelegation = \{(port, se, seport) \mid port \in ea.type.ports \wedge se \in ea.subElements \wedge$<br>$se \in ea.subElements \wedge seport \in se.type.ports$<br>$\wedge$ (<br>$(\exists lb \in ea.localBindings \wedge IsSource(lb, ea, port) \wedge$<br>$IsDestination(lb, se, seport))$<br>$\vee$<br>$(\exists lb \in ea.localBindings \wedge IsSource(lb, se, seport) \wedge$<br>$IsDestination(lb, ea, port))$<br>$\vee$<br>$(\exists rb \in ea.remoteBindings \wedge IsEndpoint(lb, se, seport) \wedge$<br>$IsEndpoint(lb, ea, port))$<br>$)\}$<br>$PropagatedPorts = \dots$<br>$FixedPorts = \dots$<br>$RealizedFeatures = \{f \mid f \in ea.features \wedge \neg IsFeatureMapping(f.value)\}$<br>$AttributesInFeatures = \dots$<br>$DelegatedFeatures = \{(f, se, sef) \mid se \in ea.subElements \wedge f \in ea.features \wedge$<br>$sef \in se.features \wedge IsFeatureMapping(f.value) \wedge$<br>$f.value.featureName = sef.featureName \wedge$<br>$f.value.subElement = se\}$ |
|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

Fig. 21 Auxiliary sets for composite element architecture

```

 $\forall e \in \text{Elements} : \text{ElementHasEA}(e, \text{SocketFactorySkeleton}) \Rightarrow ($ 
  AvailablePorts( $e, \{Mw, \text{CallOut}\}$ )  $\wedge$ 
  SubElements( $e, \{\text{SocketFactory}, \text{Skeleton}\}$ )  $\wedge$ 
  SupportedFeatures( $e, \{\text{Security}\}$ )  $\wedge$ 
  RequiredAttributes( $e, \emptyset$ )  $\wedge$ 

  -- definition of sub-elements
  -- SocketFactory sub-element
  AllowedEAsForSubElement( $e, \text{SocketFactory}, \{\text{SocketFactoryProvider}\}$ )  $\wedge$ 
  PropagateDeployment( $e, \text{SocketFactory}$ )  $\wedge$ 
  AvailablePorts( $e, \text{SocketFactory}, \{\text{Factory}\}$ )  $\wedge$ 
  -- Skeleton sub-element
  AllowedEAsForSubElement( $e, \text{Skeleton}, \{\text{RMISkeleton}\}$ )  $\wedge$ 
  PropagateDeployment( $e, \text{Skeleton}$ )  $\wedge$ 
  AvailablePorts( $e, \text{Skeleton}, \{\text{CallOut}, Mw, \text{Factory}\}$ )  $\wedge$ 

  -- port bindings
  PortBinding( $e, \text{Skeleton}, \text{Factory}, \text{SocketFactory}, \text{Factory}$ )  $\wedge$ 
  -- port delegation
  PortDelegation( $e, \text{CallOut}, \text{Skeleton}, \text{CallOut}$ )  $\wedge$ 
  PortDelegation( $e, Mw, \text{Skeleton}, Mw$ )  $\wedge$ 

  -- signature propagation
  -- enforcing fixed signatures
  -- enforcing values of the realized features
  -- defining attributes referred in feature values
  -- feature delegation
  FeatureDelegation( $e, \text{Security}, \text{SocketFactory}, \text{Security}$ )
 $)$ 

```

Fig. 22 Example: Image of SocketFactorySkeleton composite element architecture

element  $se$ , sub-element feature  $sef$ ) such that  $f$  is mapped to  $sef$ . Mathematically, the set and relations can be defined by the expressions in Fig. 21.

5.2.4 Example: Image of SocketFactorySkeleton. Referring back to the example from Sections 2 and 4, consider the actual parameters  $ea = \text{SecureSkeleton}$  and  $eas = \{\text{SerializedServerSkeleton}, \text{SocketFactorySkeleton}, \text{SocketFactoryProvider}, \text{RMISkeleton}, \dots\}$ ; the transformation will produce the image presented in Fig. 22.

### 5.3 Transforming specification of a Distribution Architecture

To illustrate the transformation, we provide an example of the *RPC* distribution architecture (Fig. 23); a more detailed description of the template, as well as a description of the formal parameters and auxiliary sets, is given in Appendix B.

5.3.1 Example: Image of *RPC* distribution architecture. Similar to element architecture, the image of *RPC* distribution architecture (Fig. 23) formulates constraints in the form of an implication. In this case, however, the constraints concern the whole connector, i.e., the image formulates the constraints on the current connector employing the distribution architecture *RPC*, the specification of which was subject to the transformation. As an aside, the element of the *Connector* set from the meta-model, representing the connector, is not explicitly mentioned (in contrast to  $e$  in the case of elements), since we assume it is unique. This assumption is correct, as a single connector theory represents a single connector (represented by the unique element of the *Connector* set).

Similar to element architectures, the connector can feature only the roles defined by the communication style refined by the *RPC* distribution architecture (the *role* relation is constrained by *AvailableRoles*). Further, the cardinality of the roles is reflected (via constraining *roleInstance* by *RolesWithSingleCardinality* and

```

EmployedDA(RPC)  $\Rightarrow$  (
  AvailableRoles({Caller, Callee})  $\wedge$ 
  RolesWithSingleCardinality({Callee})  $\wedge$ 
  RolesWithMultipleCardinality({Caller})  $\wedge$ 

  AvailableUnits({For_Caller, For_Callee})  $\wedge$ 
  UnitsWithSingleCardinality({For_Callee})  $\wedge$ 
  UnitsWithMultipleCardinality({For_Caller})  $\wedge$ 

  SupportedFeatures({Security})  $\wedge$ 

  -- definition of units
  -- For_Caller unit
  AllowedEAsForUnit(For_Caller, {LoggedClientStub, RMISStub})  $\wedge$ 
  AvailableUnitPorts(For_Caller, {CallIn, Mw})  $\wedge$ 
  -- For_Callee unit
  AllowedEAsForUnit(For_Callee, {SerializedServerSkeleton,
                                  SocketFactorySkeleton})  $\wedge$ 
  AvailableUnitPorts(For_Callee, {CallOut, Mw})  $\wedge$ 

  -- association of units' ports with roles
  PortAssociatedWithRole(For_Caller, CallIn, Caller)  $\wedge$ 
  PortAssociatedWithRole(For_Callee, CallOut, Callee)  $\wedge$ 

  -- remote port bindings
  RemoteBinding(For_Callee, Mw, For_Caller, Mw)  $\wedge$ 

  -- feature delegation
  FeatureDelegation(Security, For_Callee, Security)  $\wedge$ 
)

```

Fig. 23 Example: Image of RPC distribution architecture

*RolesWithMultipleCardinality*, respectively). Similarly, only the units defined by the RPC distribution architecture are allowed (*unit* is constrained by *AvailableUnits*), and their cardinality is reflected (via constraining *unit* by *UnitsWithSingleCardinality* and *UnitsWithMultipleCardinality*, respectively). As for the further constraints, the connector may refer only to the features realized by *RPC* (*featureRequirements* is constrained by *SupportedFeatures*). Large part of the template focuses on properties of the units and their bindings. For each unit, the set of element architectures that can be employed for its refinement is explicitly prescribed (constraining the *unit* and *elementArchitecture* relations by *AllowedEAsForUnit*). This ensures the refinement consistency of the units. Further, it is desirable to reflect that a unit can feature only the ports defined by its element type prescribed in *RPC* (constraining *unit* and *ports* by *AvailableUnitPorts*). Based on this, the association of roles with units' ports is captured (via constraining *role* and *unit* by *PortAssociatedWithRole*). The remote port bindings are reflected by forcing the signatures of the involved ports to be equal (*unit* and *ports* are constrained by *RemoteBinding*). Finally, the feature delegation has to be captured (*unit* and *features* are constrained by *FeatureDelegation*).

#### 5.4 Transforming Connector Requirements and Deployment specifications

To illustrate the transformation of connector requirements and deployment specifications, we provide an example of the *CashDesk\_to\_Inventory* connector (Fig. 24); the template, as well as a description of the formal parameters and auxiliary sets, is given in Appendix B.

5.3.3. Example: Image of *CashDesk\_to\_Inventory* connector requirements and deployment. The image formulates constraints on the selection of predefined distribution and element architectures with respect to required features and

```

DefinedEndpoints({CashdeskCaller, InventoyCallee}) ∧
AllowedDAs( {RPC, LocalProcedureCall, ...} ) ∧

-- definition of connector endpoints
-- CashdeskCaller endpoint
HasRole(CashdeskCallerInstance, Caller) ∧
HasSignature(CashdeskCaller, InventoryItf) ∧
IsDeployedOn(CashdeskCaller, DockA) ∧
-- InventoyCallee endpoint
HasRole(InventoyCallee, Callee) ∧
HasSignature(InventoyCallee, InventoryItf) ∧
IsDeployedOn(InventoyCallee, DockB) ∧

-- definition of endpoints' features
EndpointFeatureRequirements(CashdeskCaller, Logging, LoggingToFileFoo) ∧

-- definition of global connector features
ConnectorFeatureRequirements( Security, SSL)

```

Fig. 24 Example: Image of CashDesk\_to\_Inventory connector requirements

deployment. Since communication style is the binding concept between connector requirements and predefined artifacts, the constraints are focused on the actual endpoints of the connector (i.e., actual role instances).

Therefore, it is first necessary to define the available endpoints of the connector (via constraining the *units* and *roleInstances* relations by *DefinedEndpoints*). Further, the set of distribution architectures that can be employed for refinement of the connector is explicitly prescribed (*distributionArchitecture* is constrained by *AllowedDAs*). For each endpoint, its role is explicitly defined (*role* is constrained by *HasRole*), the signature of the associated component is assigned to the endpoint (*signature* is constrained by *HasSignature*), and its required deployment is enforced (via constraining *roleInstances* and *node* by *IsDeployedOn*). Further, each feature requirement imposed directly on a particular endpoint is reflected (*roleInstances* and *features* are constrained by *EndpointFeatureRequirements*). Finally, the required values of the global connector features are enforced (via constraining *connectorFeatures* by *ConnectorFeatureRequirements*).

## 6. ARCAS IN ALLOY

In this section, after providing a very brief introduction to the Alloy modeling language (while concentrating only on necessary concepts), we show how a connector theory can be expressed in Alloy. We also discuss the approaches for selecting an optimal CIC with the help of Alloy Analyzer.

### 6.1 Introduction to Alloy

This section gives a brief introduction to the Alloy modeling language - a formal modeling language based on a first-order predicate logic with operators from set theory (e.g., intersection, cartesian product), a relational algebra (e.g., relational join, transitive closure), and a basic arithmetic (e.g., integer operations, set cardinality) [J02, J06]. The language is based on the notions of *signature* and *relation*. A signature is a set of abstract elements; relations are defined upon such sets. Alloy allows constraining relations by first-order logic formulas called *facts*. A fact can employ named predicates and function symbols. In general, an Alloy specification represents a first-order logic theory (*Alloy theory*) determined by the signature, relation, and facts definitions.

The Alloy Analyzer, the associated model solver, can either find a model of an Alloy theory<sup>2</sup> or check all models against a given property (expressed as a fact). The Alloy Analyzer converts the given Alloy theory to a SAT formula, uses an underlying general-purpose SAT solver to resolve the formula, and then maps the found SAT solution to an Alloy theory model. As an aside, the Alloy Analyzer requires the domains of all signatures and relations to be explicitly bounded (due to the mapping to SAT).

Syntactically, a signature is defined using the `sig` keyword followed by the name of the signature and the list of fields defining relations. The field definitions can contain additional constraints related to the associated relation. Fig. 25 shows a part of a signature definition.

```

01  one sig Connector {
02    distributionArchitecture: one DistributionArchitecture,
03    units: UnitName one -> set Unit,
04    connectorFeatures: FeatureName set -> lone FeatureValue
05  }
06
07  abstract sig S_JavaInterface extends S {}
08  one sig FactorySignature extends S_JavaInterface {}
09
10  abstract sig Element {
11    subElements: SubElementName lone -> lone SubElement,
12    ...
13  }
14  sig SubElement extends Element {}

```

Fig. 25 Example of signature definition in Alloy

Obviously, the syntax complies with the object-oriented paradigm by defining signature as a structure constituting fields. Moreover, the Alloy syntax allows signature nesting (the `extends` construct, lines 8, 9) which resembles subtyping, and also defines an abstract signature (line 8) akin to an abstract super-type definition.

Semantically, this example determines the set `Connector` being in the `distributionArchitecture`, `units`, and `connectorFeatures` relations. The `Connector` set has to have exactly one element (i.e., a *singleton* signature). As an aside, `lone` denotes a set with at most one element, while `some` denotes a set with at least one element. The declaration of `distributionArchitecture` introduces a relation between the sets `Connector` and `DistributionArchitecture`, where for each element of `Connector` there is exactly one element of `DistributionArchitecture` in the relation `distributionArchitecture`. The declaration of `units` introduces a relation between three sets `Connector`, `UnitName` and `Unit`. For each `c` of `Connector` and each `un` of `UnitName` there is a set of `u` of `Unit`, so that the triple  $\langle c, un, u \rangle$  is in `units` (denoted by `-> set`), and also for each `c` and `u` there is at exactly one `un` so that  $\langle c, un, u \rangle$  is in `units` (denoted by `one ->`). The relation `connectorFeatures` is defined analogously.

In general, a signature `S` is interpreted as a relational algebra expression in such a way that each its field `F` represents a relation between `S` and the signatures introduced by `F`; in a similar way, nesting of signatures determines a subset relation on the signatures; specifically, an abstract signature contains elements of its nested signatures only.

An Alloy fact expresses a constraint over the sets and relations introduced by signature declarations. Each fact in a specification is an axiom of the theory determined by the Alloy specification. The following example shows a fragment of a fact definition.

<sup>2</sup> In the Alloy documentation, Alloy theory is called Alloy model and a model of an Alloy theory is called Alloy model instance.

```

01 pred isSubElement[parent: Element, child: SubElement] {
02     child in univ.(parent.subElements)
03 }
04
05 fact ElementHierarchyIsTree {
06     -- There are no cycles among elements
07     no iden & ^{ parent: Element, child: SubElement |
08         isSubElement[parent, child] }
09
10     -- Every element which is not a unit has exactly one owner
11     all e: Element | e in Element - Unit <=>
12         one owner: Element | isSubElement[owner, e]
13 }

```

The fact `ElementHierarchyIsTree` (lines 5-13) describes properties of the `subElements` relation using the named predicate `isSubElement`. In principle, the fact expresses that there are no cycles among the elements of `FullFledgedEA` with respect to the relation `subElements`, and that each element of `FullFledgedEA` has exactly one owner with respect to `subElements`. The fact consists of two clauses (lines 7-8, 11-12) bound by conjunction (expressed implicitly by a new line, or explicitly by `&&`).

The first clause can serve us for explaining the basic Alloy constructs related to facts. The identity relation (`iden`) has an empty intersection with the transitive closure of the relation defined in the curly brackets (`&` stands for set/relation intersection, `^` for transitive closure of a relation, and `no` denotes that the resulting set/relation is empty.). The curly brackets define a relation as follows: the relation contains all the pairs of elements `parent` and `child` of `FullFledgedEA`, such that they satisfy the predicate `isSubElement` (lines 1-3). This predicate is satisfied, if and only if `child` is in the relation `subElements` with `parent` and an arbitrary element of `EASE`. In other words, the tuple  $(parent, x, child)$  is in the relation `subElements` for an  $x$  in `EASE`. The operator `|` is read as “such that”. The `in` operator stands for set/relation inclusion. The operator dot (`.`), as in `parent.subElements`, denotes relational join over a relation (`subElements`) using a joined relation (in this case an unary relation `parent` containing a single element of `FullFledgedEA`). This relational join results in a new relation (in this case a relation of two sets: `EASE` and `FullFledgedEA`). In the example, the first join is followed by an outer join, where the joined relation is the entire domain of `EASE` denoted by `univ`. Consequently, the outer join yields a set of the right-most elements in `subElements`. Since the operator `[]` stands for relational join as well, the double join here can be rewritten also as `parent.subElements[univ]` or even `subElements[parent][univ]`. Note that `[]` also indicates the arguments of a predicate (e.g., `isSubElement`), the interpretation of the operator depends on a particular context.

## 6.2 Connector Theory in Alloy

In this section, we describe how a connector theory (Section 5) can be represented by means of an Alloy model. We focus on representing the connector theory meta-model (**Error! Reference source not found.**) in Alloy, describe a realization of its sets, and provide examples of how the connector-theory constraints (i.e., the images of element architectures, distribution architectures and requirements) can be reflected in Alloy.

**6.2.1 Representing the Meta-model.** In order to express a connector theory in Alloy, first it is necessary to express this way its meta-model from Fig. 13. Since the meta-model is based on sets and relations, it can be represented in Alloy in a straightforward way as illustrated below. For purpose of the following examples we have modified the relations `features`, `attributes`, and `ports` so that their second and third fields were switched (e.g., the `features` relation was changed from  $Element \times Signature \times Port$  to  $Element \times Port \times Signature$ ). The rationale is a more convenient Alloy syntax for expressing the constraints in the modified case. Note that

the definition of the sets and relations includes also several constraints concerning cardinality of the relations.

```

abstract sig DeploymentDock {}
-- similar for the following sigs

-- Role                               ElementArchitecture
-- DistributionArchitecture           SubElementName
-- UnitName                           Attribute
-- Port                               AttributeValue
-- LocalProvidedPort                 Feature
-- LocalRequiredPort                 FeatureValue
-- RemotePort                         Signature

abstract sig RoleInstance {
  signature: one Signature,
  role: one Role
}

abstract sig Element {
  elementArchitecture: one ElementArchitecture,
  ports: PortName set -> lone Signature,
  subElements: SubElementName lone -> lone SubElement,
  dock: one DeploymentDock,
  features: FeatureName set -> lone FeatureValue,
  attributes: AttributeName set -> lone AttributeValue
}

sig Unit extends Element {
  roleInstances: RoleInstance
}

sig SubElement extends Element {}

one sig Connector {
  distributionArchitecture: one DistributionArchitecture,
  units: UnitName lone -> set Unit,
  connectorFeatures: Feature set -> lone FeatureValue
}

```

The definition of the connector theory meta-model includes the integrity constraint `ElementHierarchyIsTree` articulated in Section 0.

6.2.2 Realizing the Meta-model Sets. Realization of the fixed sets by enumerating all their elements is an inherent part of a connector theory. The individual elements are introduced by the given specifications. In Alloy, each such element is reflected by creating a singleton set. As an example, consider realization of the `DeploymentDock` set. For each dock in the deployment specification, a singleton set will be created. Thus, for the deployment specification from Section 4.2.3, the following Alloy code will be created:

```

one sig DockA extends DeploymentDock {}
one sig DockB extends DeploymentDock {}

```

In the case of `FeatureValues`, the representation is more complicated, since feature values may be hierarchical, introducing feature value parameters. In Alloy, we represent a parameter  $p$  by a designated relation  $p$  between two `FeatureValues`.

Moreover, in the context of an element architecture, the value of a parameter may be represented by an attribute and expected to be given later in the requirements specification (e.g., `fileName` in `LoggingToFile` in Section 4.1). This is expressed in Alloy by marking the associated signature by `abstract` (instead of `one`). This means that multiple concrete feature values are expected to be explicitly defined (each represented by a singleton subset), while the current signature defines their required

structure. As an example, consider the `LoggingToFile` feature value defined in the `FileLogger` element architecture, which includes the `fileName` parameter as its attribute.

```
abstract sig LoggingToFile extends FeatureValue {
  fileName: one FeatureValue -- feature value parameter
}
```

A concrete value explicitly defines the value of the parameter. As an example, consider the `LoggingToFile_InventoryLog`, given by the requirements specification in Section 4.2, which assigns to `fileName` the “`inventory.log`” value (represented by `InventoryLog`).

```
one sig InventoryLog extends FeatureValue {} -- stands for “inventory.log”
one sig LoggingToFile_InventoryLog extends LoggingToFile {} {
  fileName = InventoryLog -- assignment of a particular parameter value
}
```

As for the alterable sets, in case of the `Element` set the individual elements are created automatically by the Alloy Analyzer constraint solver according to the related constraints. In case of the `Signature` set, in contrast to parameterized feature values, concrete cases of the parameterized signatures are expected to be created by the Alloy Analyzer. The rationale is that while all the concrete values of the feature value parameters are explicitly given by the requirements specification, the concrete values of signature parameters depend on the actual composition and bindings of the elements, created by the Alloy Analyzer. Technically, the Alloy representation of a signature parameterized by a signature variable is similar to the representation of a feature value parameterized by an attribute; i.e. via a dedicated relation. The fact that signatures are expected to be generated by the Alloy Analyzer is reflected by not marking the associated signature abstract. As an example, consider the `RMI` signature with has the `javaSig` parameter, introduced by `RMI` skeleton in Section 4.1.

```
sig RMI extends Signature {
  javaSig: one Signature -- signature parameter
}
```

**6.2.3 Representing the Constraints.** In this section, we illustrate the representation of a connector theory (the predicates in particular) in Alloy using examples of constraints and predicates for element architecture (Sections 5.1 and 5.2). First, the associated fixed sets have to be defined, including `Ports`, `SubElementNames`, `Features`, and `Attributes` sets. All the elements to be defined are obtained by scanning the specification and the associated element type. For each element architecture, a single element of the `ElementArchitecture` set is created, representing the element architecture itself. Thus, for the `FileLogger` element architecture specification from the running example (Section 4.1), the following Alloy code will be created:

```
-- Ports
one sig CallIn extends LocalProvidedPort {}
one sig CallOut extends LocalRequiredPort {}
-- SubElementNames
-- Features
one sig Logging extends Feature {}
abstract sig LoggingToFile extends FeatureValue {
  fileName: one AttributeValue
}
-- Attributes
one sig FileName extends Attribute {}
-- Element architecture
one sig FileLogger extends ElementArchitecture {}
```

The constraint itself is a direct representation of the clause of the connector theory illustrated in Fig. 18. In Alloy, a clause is represented by a fact. The fact syntax in

Alloy is very similar to the syntax introduced in Fig. 18, slightly differing only in syntax of logical operators (also the conjunction after each predicate is in Alloy represented by a new line), sets, and application of predicates. Thus, for the `FileLogger` element architecture specification, the following Alloy fact will be created:

```
fact FileLogger {
  for all e: Element | ElementHasEA[e,FileLogger] => {
    AllowedDeployment[e, DockA]
    AvailablePorts[e, CallIn + CallOut]
    SubElements[e, none]
    SupportedFeatures[e, Logging]
    RequiredAttributes[e, FileName]

    -- signature propagation
    PropagateSignature[e,CallIn,CallOut]

    -- enforcing fixed signatures

    -- enforcing values of the realized features
    FeatureHasValue[e, Logging, LoggingToFile]

    -- defining attributes referred in feature values
    FeatureValueUsesAttribute[e, Logging, fileName, FileName]
  }
}
```

As for predicates, in Alloy the definition of a predicate is reflected by a specific construct, which is similar to the definition of a function in a programming language (i.e., it defines the name, list of the parameters, and the actual body). As a simple example, consider the `ElementHasEA` predicate from Table 1. Note, that we have replaced the test for inclusion in a relation with relational join.

```
pred ElementHasEA[e: Element, ea: ElementArchitecture] {
  e.elementArchitecture = ea
}
```

To give a more elaborate example of a predicate definition, emphasizing the advantages of the Alloy language, we present the `FeatureValueUsesAttribute` predicate from Table 1. In this case, since a feature-value parameter is represented by a dedicated relation (e.g., `fileName` in the `LoggingToFile` feature value), we exploit the possibility of giving a relation as a parameter to an Alloy predicate. We also use relational join to express the context of the constraint (we impose constraint on a feature value and an attribute value, which are identified by their context with respect to features and attributes relation, respectively).

```
pred FeatureValueUsesAttribute[e: Element, feature: Feature,
  parameter: FeatureValue -> AttributeValue, attr: Attribute] {
  e.features[feature].parameter = e.attributes[attr]
}
```

Compared to the logical formula from Table 1 reflecting the predicate, it is clear that the Alloy representation is much more concise and comprehensive, as it resembles an expression in a regular object-oriented programming language.

In a similar vein, consider the `PortBinding` predicate from Table 2. Note, that in this case the relational join allows to express the context of the constraint transitively (as combination of relations `subElements` and `ports`).

```
pred PortBinding[e: Element,
  subelement1: SubElementName, port1: LocalRequiredPort,
  subelement2: SubElementName, port2: LocalProvidedPort] {
  e.subElements[subelement1].ports[port1]
  =
  e.subElements[subelement2].ports[port2]
```

```

}

```

Further, consider the `AvailablePorts` predicate from Table 1. It is an example of transitive constraint on a second element of a ternary relation, where the third element is not relevant.

```

pred AvailablePorts[e: Element, availablePorts: set Port] {
  e.ports.univ = availablePorts
}

```

In a similar way, the Alloy representation of the `AvailableRoles` predicate from Section 0 illustrates this case of irrelevant relation item on a transitive constraint.

```

pred AvailableRoles[roles: set Role] {
  Connector.units[univ].roleInstances.role = roles
}

```

Finally, as the most complex example of a predicate, consider the `RemoteBinding` predicate from 5.3. Note that also in this case, the use of relational join greatly reduces the complexity of the actual Alloy representation.

```

pred RemoteBinding[unit1: UnitName, p1: RemotePort,
  unit2: UnitName, p2: RemotePort] {
  Connector.units[unit1].ports[p1] = Connector.units[unit2].ports[p2]
}

```

### 6.3 Selecting an Optimal CIC

The set of models of a connector theory can be of a large cardinality; nevertheless just a single model is needed. Thus, it is necessary to make a choice and select the “best” one (the best CIC). A practical necessity is to automate the selection process.

There are different criteria for judging what the “best” CIC is, ranging from memory consumption and CPU utilization, latency and throughput, to robustness, reliability, and stability, represented by means of valuation of the elements involved in CIC. Naturally, it is necessary to find rules for composability of the criteria. This leads to an optimization problem where the task is to find an optimal CIC given a valuation of its elements (their element architectures in particular) by applying the rules. A simple example of such valuation is a manual assignment of a fixed cost to each element architecture where the valuation of CIC is the sum of costs for all the elements in it. The CIC with the lowest cost is pronounced the “best”.

The Alloy language itself does not provide any explicit support for solving optimization problems. In principle, there are three possible approaches to employing the Alloy framework for finding an optimal CIC: (i) to encode the optimization problem into a “standard” Alloy theory; (ii) to extend the Alloy language and Alloy Analyzer accordingly; and (iii) to enumerate all the models of a given connector theory and select an optimal one by applying valuation criteria outside the resolution process (an enumeration of all models is a standard feature of the Alloy Analyzer).

We do not consider other approaches, such as extending the Alloy Analyzer without changing the Alloy language (i.e., instrumenting the Analyzer with custom heuristics), or moving from Alloy to another constraint-solving technique. Preferring relative simplicity and efficiency, we have used the approach (iii) for the experiments and evaluation.

**6.3.1 Encoding Optimization Problems in Alloy.** Modeling optimization problems in Alloy is based on expressing a partial order over the set of CICs where the “best” CIC is the least element. Since the order of a particular CIC is typically determined by the employed element architectures, it is necessary to define a global partial order of all the available element architectures, as well as a way of inferring the CIC partial order from this global order.

Defining such an inferring is a challenge. A simple approach might be to use a fixed cost of element architectures, modeled via Alloy natural numbers, where the

aggregation is the cost sum of the employed element architectures in a particular CIC. However, the usage of numbers and arithmetic in an Alloy theory greatly increases its complexity. Another option is to employ the standard Alloy relations to express the partial order of element architectures instead of comparing cost values. However, it would be necessary to assess the order of composite element architectures based on the order of its sub-element architectures, which is also a challenge. Another option is also to employ the state-of-the-art alternatives for solving Alloy theories [EGT11, EGT10], that offer better support for arithmetic and unbounded integer types.

**6.3.2 Extending the Alloy Framework.** The Alloy language could be extended to allow for specifying optimization criteria. The Alloy Analyzer would have to translate such optimization criteria into a SAT formula. Specifically, such an extension can be achieved by employing pseudo-boolean (PB) formulas instead of SAT formulas. Here, an assumption is that the Alloy Analyzer supports a PB solver such as MiniSAT+ [ES06] (this appears to be realistic in the future, since the MiniSAT solver is already supported by Alloy Analyzer). Overall, such an extension of the Alloy framework requires a major modification of its current implementation.

**6.3.3 Enumerating All Connector Theory Models.** A standard feature of the Alloy Analyzer is the enumeration of all models of a given Alloy theory by incremental execution of its underlying SAT solver. This feature can be effectively exploited for finding an optimal CIC. Here, the Alloy Analyzer is let to find all the models of a given connector theory and a valuation of the models is provided by a dedicated external tool (i.e., a partial order upon the models is created); finally an optimal CIC is determined based on the partial order.

By delegating the optimization to a dedicated external tool, this method does not require expressing optimization criteria in CT. Consequently, the model valuation strategy is not limited by the expressive power of the Alloy language, but it is entirely determined by the options the external tool provides.

Even though it is necessary to explore all the models, this method is still tangible, since each of the models is evaluated separately.

## 7. RELATED WORK

In general, the related work spans three areas: (i) composing software with the help of constraint solving, (ii) Alloy-based resolution and verification of component architectures, and (iii) automated connector synthesis.

(i) Composing software with the help of constraint solving. Constraint solving techniques have been already used for the variety of tasks in software composition stretching from automated dependency management [LBP08] to verification of composition in product lines [TBKC07]. In [LBP08] the idea is to employ a SAT solver to resolve dependencies; this approach was used in several contemporary software tools such as OSGi implementation Equinox p2 and Maven.

The approach presented in [TBKC07] is based on formalizing the notion of a feature model by introducing a specific feature algebra. Having such formalization available, a SAT solver is employed for verification of safe feature-model composition in a product line.

Compared to ARCAS, both methods leverage on simple propositional formulas for capturing the given problem. Such formulas merely express variability and transitivity, but do not reflect more complex properties of the software parts to be composed (e.g., interface bindings, runtime environment requirements, and features).

(ii) Alloy-based resolution and verification of component architectures. Alloy has been already extensively used in the domain of CBSE for the purpose of both property checking and model finding. In [GMK02], the authors examine the feasibility of using architectural constraints as the basis for specification, design, and implementation of self-organizing architectures in Darwin. In this context, Alloy

serves as a tool for an automated resolution of the self-organizing reconfigurations from an inconsistent to a consistent state in terms of a particular architectural style.

In a similar way, [HI10] employs Alloy for specification of the possible architecture reconfiguration actions in the context of a generic component model based on OSGi. Alloy is employed in two ways: (a) architecture change verification and (b) architectural change planning. The former focuses on soundness of the reconfiguration actions and preservation of the properties specific to a particular architectural style. The latter comprises finding a fitting sequence of reconfiguration actions from the current consistent state to a given consistent state while preserving a particular architectural style in all intermediate states.

A methodology for verifying soundness of self-configuration scenarios via their specification in Alloy is presented in [TMS10]. The underlying formal method – FracToy – formalizes a concrete self-configurable system, as well as the corresponding self-organizing actions. The verification targets both static and dynamic properties. Compared to our approach, models of the Alloy theory only prove the consistency/soundness of the theory and are not used for any other purpose.

In [KG06], Alloy is employed for formal specification of various architectural styles. The main goal is to check the important properties of an architectural style such as consistency, satisfaction of a predicate over an architectural style instance, composability, and refinement of architectural styles. The Alloy formalization of an architectural style is obtained programmatically from its Acme specification (which serves a similar purpose as our CDL specification).

Apart from reconfigurations, in [JS00] a formal specification of valid component compositions in the COM component model is analyzed, while in [MS08] a fully-fledged Alloy formalization of the Fractal component model is presented.

(iii) Automated connector synthesis. The ARCAS method stems from our previous research [BP04, BB05, B06, BHP06]. In [BP04, B06], a connector model designed for automated connector synthesis based on a high-level specification is presented. Despite leveraging on similar concepts to those presented in Section 2, it does not explicitly capture NFPs. The key difference lies in the way a connector configuration is selected - it is performed via term matching in Prolog. Breaking the separation of abstraction levels, the method enforces inclusion of Prolog terms directly in the requirements specification.

An extensive effort has been put into research of automated synthesis of connectors ensuring application-layer and middleware-layer interoperability [IBB11]. For brevity, we call the former API mediation and the latter middleware bridging. While we have focused on cases where a component/service being deployed is to be connected to another one (potentially already deployed and running), [IBB11] exploits a slightly different scenario – connecting solely the already deployed and running components/services. The connected components may require both API mediation and middleware bridging. Thus, letting NFPs aside, in [IBB11] the input is an API specification and a determination of a particular middleware for each of the connected components; the goal is to synthesize (in an automated way) a connector implementation which does the mediation and/or bridging. For comparison, in ARCAS the input is a deployment and requirements specification; the goal is to find a fitting connector implementation employing a suitable middleware/middleware bridging (a potential API mediation is expected to be done by a separate dedicated component).

In [RRSGWT05], the concept of connector is introduced for CCM. The goal is to benefit from a light communication middleware – lighter than CORBA – in CCM-based applications by introducing connectors (strictly separating the component communication from business logic). The connector-generation process employs an extended IDL for defining connector templates. These are however not composable. Similar to ARCAS, connector specification is accompanied by a

description of connector-specific features in an extended OMG D&C specification. Connectors are generated by manually selecting a particular template parameterized by the actual interfaces; the connector templates have also to be created manually.

In [RCGT09], targeting model-based generation of component-based applications from UML-MARTE models, the concept of compositional connectors is also employed. A connector's architecture has to be explicitly captured (manually) by an UML-MARTE model. The abstractions describing a connector's structure are similar to those presented in Section 2.

## 8. DISCUSSION AND EVALUATION

In this section, we discuss the important decisions and tradeoffs we had to take, and also mention open issues we still face. In particular, these include (i) interpretation of a connector theory model, (ii) focus of a connector theory, (iii) addressing NFPs, (iv) weak points of the connector abstractions, (v) finding of an optimal connector theory model, (vi) experience and case study, and (vii) moving ARCAS to other domains.

(i) *Model Interpretation.* The ARCAS method exploits the Alloy Analyzer's capability of finding a model of a given Alloy theory. In contrast to a typical usage of this feature – interpreting the existence of a model as consistency/soundness of the theory, and using a model as feedback during theory development, e.g., [ABGR08, TMS10] – the ARCAS method directly interprets/employs the found model as a CIC.

(ii) *Focus of the Theory.* The proposed connector theory differs from similar formal models [MS08] in its focus. While in [MS08] the target is captured at meta-model level, the connector theory is built at that level only partially (specifically this refers to the meta-model, Section 5). Most of the theory is built at model level, i.e., it specifies semantics of a particular connector with respect to concrete instances of abstract entities (e.g., `FileLogger` as an instance of element architecture). This is facilitated by the automated process of converting connector specification into a connector theory, whereas formal models at meta-model level are typically created manually. The ARCAS method thus brings the option of reasoning about a particular connector instead of reasoning solely about properties specific to all possible connectors.

(iii) *Addressing NFPs.* The ARCAS method makes it possible to address NFPs in the synthesized connector. This is in contrast to some of the state-of-the-art methods for automated synthesis of middleware-based connectors [IBB11]. Since the Alloy theory (Section 5) resulting from a connector specification comprises also element composability with respect to NFPs, it is possible to programmatically synthesize a connector complying with the NFP requirements. Because the Alloy Analyzer lacks specialized support for arithmetic operators, it is feasible to address only qualitative NFPs (i.e., those based on enumeration of values), whereas efficient addressing of quantitative NFPs (such as latency) would be hard to achieve. However, there are constraint-solving techniques providing advanced support for arithmetic operations and working with inequalities [DMB08], which might be sufficient for addressing quantitative NFPs. In this respect, a challenging issue is to integrate these constraint-solving techniques with the Alloy framework, or adapt the ARCAS method into another framework based on such techniques.

(iv) *Issues not addressed.* The presented connector concepts do not reflect (a) cardinality of ports and sub-elements (e.g., important when each of the multiple client stubs has to be served by a dedicated server skeleton), and (b) compound port signatures (a typical phenomenon in cases where a server unit supports a number of middleware protocols simultaneously, such as RMI and SOAP). In this paper, these concepts were left out for simplicity, but the ARCAS method can be enhanced this way. Inclusion of (a) comprises an extension to the element type and element

architecture abstract syntax, and straightforward modifications of the transformation of element architectures, as well as distribution architectures. Inclusion of (b) involves modifications of the connector theory and the associated transformations. As an aside, both (a) and (b) have been already experimentally researched.

No formal-method-based attention is given to an explicit behavior specification and behavior matching of element architectures. This is because the middleware-based behavior of connectors is rather simple and driven by the communication style. Thus the behavior of an element architecture is to be driven merely by the element type it implements; therefore, we adopted a name convention to express the expected behavior of element types. Thus the responsibility of the behavioral compliance between element types and element architectures is left on the element designer (situation similar to the responsibility of correctly implementing an interface in a class).

(v) Finding of an optimal connector theory model. As presented in Section 6.3, we have employed the method of enumerating and independently valuating all the connector theory models in finding an optimal one. An open issue is to experimentally evaluate and assess feasibility of the other two methods considered in Section 6.3 (and possibly propose new ones). With respect to state explosion, a preliminary assumption is that the method based on extending the Alloy framework would be the most fitting for smaller distribution architectures (in terms of number of employed elements). For large distribution architectures, we considered as better the method based on enumeration of all connector theory models because of its low space complexity, even though it is probably of the worst time complexity. Overall, the method based on encoding an optimization problem in a “standard” Alloy theory is a trade-off between time complexity and implementation effort on one hand, and space complexity on the other.

(vi) Moving ARCAS to other domains. Although ARCAS method has been presented with the main use-case of automated resolution of middleware-based connectors, the basic idea of the method is generic and applicable in other domains. ARCAS universally applies to cases where a series of activities, that together mediate and transform an input source data or event to a number of output data or events, is to be assembled. This includes architectural patterns such as pipe-and-filter and workflows in general (including nested workflows). A typical example illustrating such a potentially nested workflow is a media player application. Such application employs a number of codecs (audio and video), filters, muxers and demuxers, which have to be correctly organized in its architecture in order to process the content from an input stored in a file or available on-line. From the ARCAS perspective, the whole architecture of a media player application can be likened to a connector, and each of the codecs, filters, muxers, and demuxers can be likened to a connector element. The key property of ARCAS, applicable in this scenario, is the automated composition of the elements while ensuring compatibility among mutually connected elements. Moreover, ARCAS method allows for imposing constraints on the whole connector in terms of required features. By adapting these concepts to the case of a media player, it is possible to declaratively enforce for instance the use of a de-interlacing filter or select a tradeoff between computational complexity and video image quality. This possibility of specifying features, which adjust the behavior and characteristics of the whole connector, is also one of the main advantages over other existing approaches (e.g., [LBP08, TBKC07]).

In addition to workflow architectures (pipe-and-filter style being an special subclass of them), ARCAS method also allows for automated construction of layered architectures, if the mutual requirements in-between the layers can be sufficiently described. An example illustrating such application in the domain of embedded real-time system may be an automated selection of hardware sensors along with corresponding device drivers and proper operating system abstraction layer. Using

ARCAS terminology, the whole software stack resembles a connector, while each sensor, device driver, and particular abstraction layer can be likened to an element. Similar to the previous example, it is possible to take advantage of the feature specification in order to for instance specify required sampling rate of a sensor.

(vii) Experience and case study. As a proof-of-the-concept, we have developed an experimental database of connector artifacts (including both the specifications and their Alloy images). The presented examples were based on simplified version of the artifacts in this database. We have employed this database in a case study involving a non-trivial part of a real-life component-based application based on the procedure-call communication style [HKW08]. Various client-server connection scenarios differing in NFPs and deployment were considered. This helped to clearly demonstrate the soundness and feasibility of the ARCAS method on a realistic example (especially feasibility of Alloy-based CIC resolution).

We have also performed several benchmarks in order to assess the performance scaling factors of ARCAS. For this purpose, the actual database of connector artifacts was generated in an automated way by renaming and by introducing new variants in the original database. Nevertheless, this approach well captures the general case, since the performance of Alloy Analyzer strongly relies on the cardinality of the sets and relations rather than on the complexity and variability of the constraints. Even though the computational complexity is exponential in principle (Alloy Analyzer employs a SAT solver), based on the measurements ARCAS is feasible to hundreds of element architectures and tens of distribution architectures. Specifically, for 100 element architectures and 10 distribution architectures the execution times are in the order of 5 seconds (MiniSAT, 512MB, Intel i5 2.6 GHz); for 200 element architectures and 20 distribution architectures the execution time is around 14 seconds. The performance can be further improved by representing connector theory directly in Kodkod [T09] (the underlying relational solver) instead of in Alloy.

Note that the size of the database actually contains only the elements architectures and distribution architectures, which are applicable with respect to the deployment and selected communication style (Section 3), thus an actual artifact database will be typically even larger.

We have also developed an EMF<sup>3</sup>-based demonstrator of a tool for automated transformations of the specifications.

## 9. CONCLUSION AND FUTURE WORK

In this paper, we presented a method for automated resolution of connector architectures based on constraint solving – the ARCAS method – as the first step in the automated generation of middleware-based connectors. An important benefit of ARCAS is the ability to address the required NFPs, which, as well as transparent distribution, is the major concern of middleware-based connectors. For purpose of ARCAS, we have introduced the concepts of middleware connectors based on hierarchical elements, stemming from hierarchical components. These concepts allow defining the individual parts of a connector in advance and thus facilitate reuse. In addition, such hierarchical connectors abstract the unnecessary details and therefore strongly contribute to the separation of concerns in CBSE. The key idea of ARCAS is to resolve a description of a particular connector instance (CIC) as a model of a theory based on a first-order logic and relational calculus – a connector theory. We have defined automated transformations, which convert the predefined connector artifact specification and connector requirements and deployment specifications to such a connector theory. This way, ARCAS can be employed whenever the requirements or deployment changes (even at runtime). As a proof-of-the-concept,

---

<sup>3</sup> <http://www.eclipse.org/modeling/emf/>

we described representation a connector theory in the Alloy modeling language. Moreover, by using the Alloy representation, we have showed feasibility of ARCAS on a real-life example. A demonstrator of an automated tool for ARCAS is available<sup>4</sup>. We have also discussed applicability of ARCAS in other domains and possible improvements.

As a future work, we aim, in addition to finalizing the automated ARCAS tool, at addressing the issues mentioned in Section 8, in particular compound port signatures and cardinality of elements. In the long term, we intend to focus on providing support for optimization problems defined in Alloy, as well as introducing support for quantitative NFPs (either by extending the Alloy framework or employing other constraint solver frameworks). Finally, we plan to explore the possibility of integrating related state-of-the art methods for middleware and application interoperability [IBB11, IST11, BPGG11] to achieve a resolution-based synthesis of emergent, full-fledged connectors.

## REFERENCES

- [ABGR08] K. Anastakis, B. Bordbar, G. Georg, and I. Ray, “On challenges of model transformation from UML to Alloy,” *Software & Systems Modeling*, vol. 9, no. 1, pp. 69–86, Dec. 2008.
- [B06] T. Bures, “Generating Connectors for Homogeneous and Heterogeneous Deployment”, *PhD dissertation*, Dept. of Distributed and Dependable Systems, Charles University in Prague, 2006.
- [BB04] L. Bulej, T. Bures, “Addressing Heterogeneity in OMG D&C-based Deployment”, Tech. Report No. 2004/7, Dep. of SW Engineering, Charles University, Prague, <http://d3s.mff.cuni.cz/publications/No2004>.
- [BB05] L. Bulej, T. Bureš: “Deploying Heterogeneous Applications using OMG D&C and Software Connectors”, Tech. Report No. 2005/10, Dep. of SW Engineering, Charles University, Prague, <http://d3s.mff.cuni.cz/publications/>, Nov 2005.
- [BGI07] S. Benmokhtar, N. Georgantas, and V. Issarny, “COCOA: COmposition in pervAsive computing environments with QoS support,” *Journal of Systems and Software*, vol. 80, Dec. 2007, p. 1941–1955.
- [BHP06] T. Bures, P. Hnetyka, and F. Plasil, “SOFA 2: Balancing Advanced Features in a Hierarchical Component Model”, Proc. of 4<sup>th</sup> *International Conference on Software Engineering Research, Management and Applications (SERA '06)*, IEEE Computer Society, Washington, DC, USA, 2006.
- [BP04] T. Bures and F. Plasil, “Communication style driven connector configurations,” in *Software Engineering Research and Applications*. Springer Berlin/Heidelberg, 2004.
- [BPGG11] G. Blair, M. Paolucci, P. Grace, N. Georgantas: “Interoperability in complex distributed systems.” In: Bernardo, M., Issarny, V. (eds.) *SFM 2011*. LNCS, vol. 6659, pp. 350–392. Springer, Heidelberg (2011)
- [BS07] S. Bliudze and J. Sifakis. “The algebra of connectors: structuring interaction in BIP”. In Proc. of *EMSOFT '07*. ACM, New York, NY, 11–20, 2007.
- [CCP11] J. Cubo, C. Canal, and E. Pimentel, “Context-Aware Composition and Adaptation based on Model Transformation,” *The Journal of Universal Computer Science*, vol. 17, 2011, pp. 777–806.
- [CL02] I. Crnkovic, M. Larsson: “Building Reliable Component-Based Software Systems.” Artech House, Inc., Norwood, MA, USA, 2002.
- [DMB08] L. De Moura and N. Bjørner, “Z3: An efficient SMT solver,” *Tools and Algorithms for the Construction and Analysis of Systems*, 2008, p. 337–340.
- [EGT10] A.A. El Ghazi and M. Taghdiri, “Analyzing Alloy Constraints using an SMT Solver: A Case Study,” *5th International Workshop on Automated Formal Methods (AFM)*, Edinburgh, United Kingdom: 2010.
- [EGT11] A. El Ghazi and M. Taghdiri, “Relational Reasoning via SMT Solving,” *FM 2011: Formal Methods*, 2011, p. 133–148.
- [ES06] N. Eén and N. Sörensson, “Translating pseudo-boolean constraints into SAT,” *Journal on Satisfiability, Boolean Modeling and Computation*, vol. 2, 2006, p. 1–26.
- [GMK02] I. Georgiadis, J. Magee, and J. Kramer, “Self-organising software architectures for distributed systems,” *Proceedings of the first workshop on Self-healing systems*, New York, New York, USA: ACM, 2002, p. 33–38.

<sup>4</sup> <http://d3s.mff.cuni.cz/~keznikl/research/arcas/>

- [HI10] K.M. Hansen and M. Ingstrup, “Modeling and analyzing architectural change with alloy,” *Proceedings of the 2010 ACM Symposium on Applied Computing - SAC '10*, 2010, p. 2257.
- [HKW08] S. Herold, H. Klus, Y. Welsch, et al., “CoCoME-The Common Component Modeling Example”, *The Common Component Modeling Example*, 2008, p. 16–53.
- [IBB11] V. Issarny, A. Bennaceur, and Y.D. Bromberg, “Middleware-layer Connector Synthesis: Beyond State of the Art in Middleware Interoperability,” *Formal Methods for Eternal Networked Software Systems*, M. Bernardo and V. Issarny, eds., Berlin / Heidelberg: Springer, 2011, pp. 217-255.
- [ISJB09] V. Issarny, B. Steffen, B. Jonsson, G. Blair, P. Grace, M. Kwiatkowska, R. Calinescu, P. Inverardi, M. Tivoli, A. Bertolino, and A. Sabetta, “CONNECT Challenges: Towards Emergent Connectors for Eternal Networked Systems,” *2009 14th IEEE International Conference on Engineering of Complex Computer Systems*, 2009, pp. 154-161.
- [IST11] P. Inverardi, R. Spalazzese, and M. Tivoli, “Application-layer connector synthesis,” *Formal Methods for Eternal Networked Software Systems*, 2011, p. 148–190.
- [J02] D. Jackson, “Alloy: a lightweight object modelling notation,” *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 11, 2002, p. 256–290.
- [J06] D. Jackson: *Software Abstractions: “Logic, Language, and Analysis.”* MIT Press, Cambridge, MA, USA, and London, England, 2006.
- [JS00] D. Jackson and K. Sullivan, “COM revisited: tool-assisted modelling of an architectural framework,” *ACM SIGSOFT Software Engineering Notes*, vol. 25, 2000, p. 149–158.
- [KG06] J.S. Kim and D. Garlan, “Analyzing architectural styles with alloy,” *Proceedings of the ISSTA 2006 workshop on Role of software architecture for testing and analysis - ROSATEA '06*, 2006, pp. 70-80.
- [LBP08] D. Le Berre, A. Parrain: “On SAT Technologies for Dependency Management and Beyond.” *Proceedings of 12th International Software Product Line (SPLC 2008)* vol. 2, 2008, p. 197-200.
- [MMP00] N. R. Mehta, N. Medvidovic, and S. Phadke: “Towards a taxonomy of software connectors,” in *Proceedings of the 22nd international conference on Software engineering*. ACM, 2000.
- [MPBH11] M. Malohlava, F. Plášil, T. Bureš, P. Hnetynka: “Interoperable DSL Families for Code Generation,” Tech. Report No. 2011/4, Dep. of Distributed and Dependable Systems, Charles University, Prague, <http://d3s.mff.cuni.cz/publications/>, Apr 2011
- [MS08] P. Merle and J.B. Stefani, “A formal specification of the Fractal component model in Alloy,” Research Report RR-6721, INRIA, <http://hal.inria.fr/inria-00338987/en/>, 2008.
- [NTER06] J. Nakazawa, H. Tokuda, W.K. Edwards, and U. Ramachandran, “A Bridging Framework for Universal Interoperability in Pervasive Systems,” *26th IEEE International Conference on Distributed Computing Systems (ICDCS'06)*, 2006, pp. 3-3.
- [OMG04] Object Management Group, “Deployment and Configuration of Component-based Distributed Applications Specification”, <http://www.omg.org/cgi-bin/doc?formal/06-04-02.pdf>, Feb 2004
- [RCGT09] A. Radermacher, A. Cuccuru, S. Gerard, and F. Terrier, “Generating execution infrastructures for component-oriented specifications with a model driven toolchain: a case study for MARTE’s GCM and real-time annotations,” *Proceedings of the eighth international conference on Generative programming and component engineering*, ACM, 2009, p. 127–136.
- [RRSGWT05] S. Robert, A. Radermacher, V. Seignole, S. Gérard, V. Watine, and F. Terrier, “Enhancing interaction support in the corba component model”, *From Specification to Embedded Systems Application*, 2005, p. 137–146.
- [SI10] R. Spalazzese, P. Inverardi: “Mediating Connector Patterns for Components Interoperability”. In: *Babar, M.A., Gorton, I. (eds.) ECSA 2010*. LNCS, vol. 6285, pp. 335–343. Springer, Heidelberg (2010)
- [TBKC07] S. Thaker, D. Batory, D. Kitchin, W. Cook: “Safe composition of product lines.” *Proceedings of the 6th international conference on Generative programming and component engineering*, ACM, 2007, p. 95–104.
- [TMD10] R.N. Taylor, N. Medvidovic, E.M. Dashofy: “Software architecture: foundations, theory, and practice.” Wiley, Hoboken, 2010.
- [TMS10] A. Tiberghien, P. Merle, and L. Seinturier: “Specifying Self-configurable Component-Based Systems with FracToy,” *Abstract State Machines, Alloy, B and Z*, vol. 5977, 2010, p. 91–104.
- [T09] E. Torlak: “A Constraint Solver for Software Engineering: Finding Models and Cores of Large Relational Specifications.” Ph.D. Thesis, MIT, February 2009

## APPENDIX

### A. ARCAS Input Specifications: Abstract Syntax




Fig. 26 Abstract syntax of Communication Style

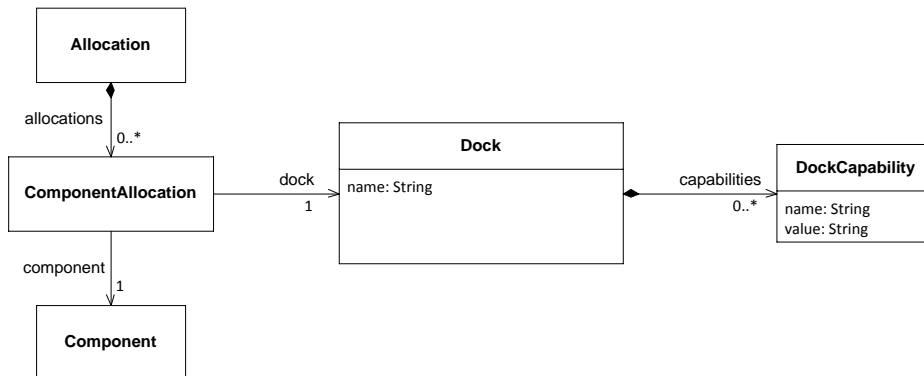


Fig. 27 Abstract syntax of Deployment Specification

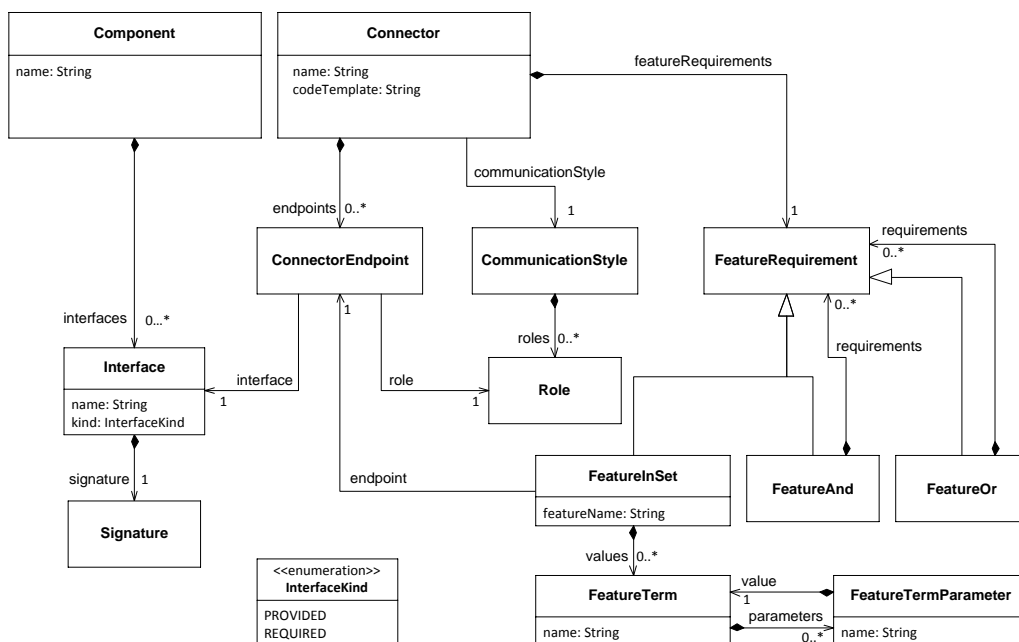


Fig. 28 Abstract syntax of Requirements Specification

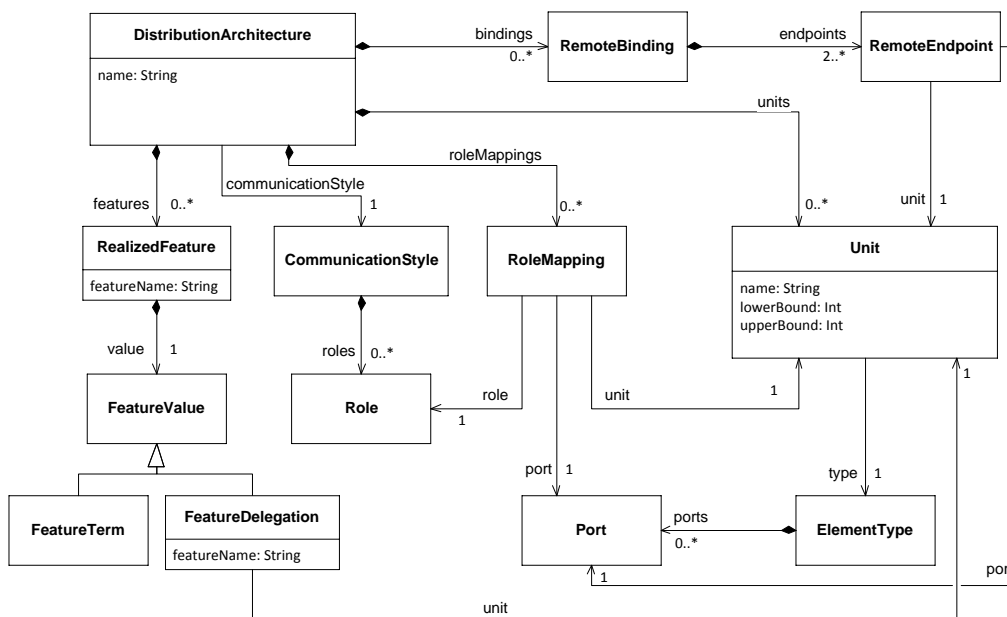


Fig. 29 Abstract syntax of Distribution Architecture specification

## B. Transformations

### B.1 Distribution architecture.

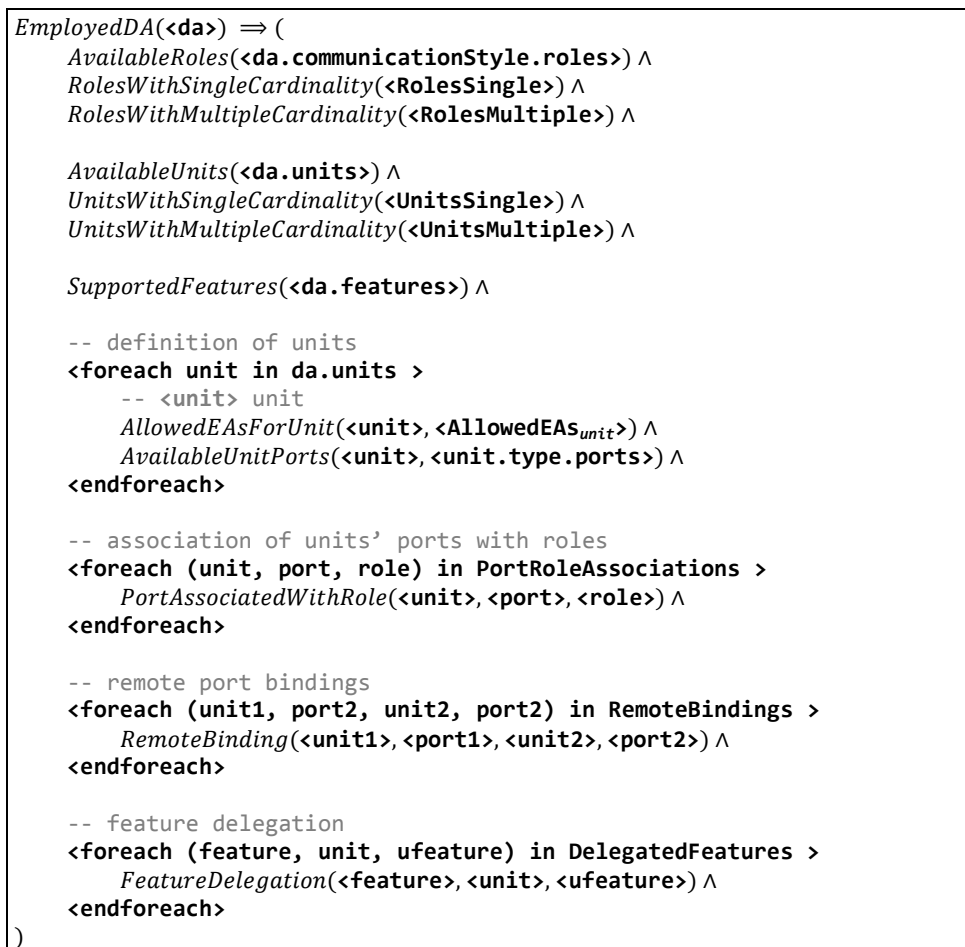


Fig. 30 Template for image of distribution architecture

```

RolesSingle= {role | role ∈ da.communicationStyle.roles ∧
              role.lowerBound = 1 ∧ role.upperBound = 1}
RolesMultiple= {role | role ∈ da.communicationStyle.roles ∧
               role.lowerBound = 0 ∧ role.upperBound = *}
UnitsSingle= {unit | unit ∈ da.units ∧
              unit.lowerBound = 1 ∧ unit.upperBound = 1}
UnitsMultiple= {unit | unit ∈ da.units ∧
               unit.lowerBound = 0 ∧ unit.upperBound = *}
AllowedEAsunit= {ea | ea ∈ eas ∧ ea.type = unit.type}
PortRoleAssociations= {(unit, port, role) | unit ∈ da.units ∧ port ∈ unit.type.ports ∧
                       role ∈ da.communicationStyle.roles ∧
                       ∃rm ∈ da.roleMappings ∧
                       rm.role = role ∧ rm.unit = unit ∧ rm.port = port}
RemoteBindings= {(unit1, port1, unit2, port2) | unit1, unit2 ∈ da.units ∧
                port1 ∈ unit1.type.ports ∧ port2 ∈ unit2.type.ports ∧
                ∃rb ∈ da.bindings ∧ e1, e2 ∈ rb.endpoints ∧
                e1.unit = unit1 ∧ e1.port = port1 ∧
                e2.unit = unit2 ∧ e2.port = port2}
DelegatedFeatures = {(f, u, uf) | u ∈ da.units ∧ f ∈ da.features ∧
                     uf ∈ u.features ∧ IsFeatureMapping(f.value) ∧
                     f.value.featureName = uf.featureName ∧
                     f.value.unit = u}

```

Fig. 31 Auxiliary sets for distribution architecture

## B2 Connector requirements and deployment.

```

DefinedEndpoints(<c.endpoints>) ∧
AllowedDAs(<AllowedDAs>) ∧

-- definition of connector endpoints
<foreach endpoint in c.endpoints >
  -- <endpoint> endpoint
  HasRole(<endpoint>, <endpoint.role>) ∧
  HasSignature(<endpoint>, <endpoint.interface.signature>) ∧
  IsDeployedOn(<endpoint>, <Deploymentendpoint>) ∧
<endforeach>

-- definition of endpoints' features
<foreach (endpoint, feature, values) in EndpointFeatureReq >
  EndpointFeatureRequirements(<endpoint>, <feature>, <values>) ∧
<endforeach>

-- definition of global connector features
<foreach (feature, values) in ConnectorFeatureReq >
  ConnectorFeatureRequirements(<feature>, <values>) ∧
<endforeach>

```

Fig. 32 Template for image of connector requirements and deployment  
(simplified – without logic operators in feature requirements)

$$\begin{aligned}
\text{AllowedDAs} &= \{da \mid da \in \text{das} \wedge da.\text{communicationStyle} = c.\text{communicationStyle}\} \\
\text{Deployment}_{\text{endpoint}} &= \{c.\text{dock} \mid c \in \text{components} \wedge \text{endpoint}.\text{interface} \in c.\text{interfaces}\} \\
\text{EndpointFeatureReq} &= \{(endpoint, feature, values) \mid \text{endpoint} \in \text{connector}.\text{endpoints} \wedge \\
&\quad \exists fr \in \text{connector}.\text{featureRequirements} \wedge \\
&\quad fr.\text{endpoint} = \text{endpoint} \wedge \\
&\quad fr.\text{featureName} = feature \wedge fr.\text{values} = values\} \\
\text{ConnectorFeatureReq} &= \{(endpoint, feature, values) \mid \text{endpoint} \in \text{connector}.\text{endpoints} \wedge \\
&\quad \exists fr \in \text{connector}.\text{featureRequirements} \wedge fr.\text{endpoint} = \emptyset \wedge \\
&\quad fr.\text{featureName} = feature \wedge fr.\text{values} = values\}
\end{aligned}$$

Fig. 33 Auxiliary sets connector requirements and deployment  
(simplified – without logic operators in feature requirements)