Traceable and information-preserving composition of architectural models

Nelis Boucké and Danny Weyns and Tom Holvoet

Report CW 538, April 2009

Katholieke Universiteit Leuven
Department of Computer Science
Celestijnenlaan 200A – B-3001 Heverlee (Belgium)

Traceable and information-preserving composition of architectural models

Nelis Boucké and Danny Weyns and Tom Holvoet

Report CW 538, April 2009

Department of Computer Science, K.U.Leuven

Abstract

C&C and deployment are base views for describing software architectures. Both views typically exist of several models that highlight particular aspects of the software. From our experience, managing the dependencies between architectural models is a challenging, yet crucial task for maintaining consistency of an architectural description. In this paper, we focus on the composition of architectural models and managing dependencies between models in such a composition. We formally define four basic relations between models and a composition operator that enables integration of several C&C models based on these relations. We elaborate on several semantic preserving properties of the composition operator, including traceability, consistency, and completeness. The relations and composition operator have been integrated in the xADL language and the AchStudio tool. We use excerpts of a distributed video-on-demand system designed with the tool as an illustrative case.

Keywords : software architecture, composition, relation, structural model, formal proof
CR Subject Classification : I.2.11, I.2.8, I.2.1

Traceable and information-preserving composition of architectural models

Nelis Boucké, Danny Weyns, and Tom Holvoet

DistriNet Labs, K.U.Leuven, Belgium
{nelis.boucke,danny.weyns,tom.holvoet}@cs.kuleuven.be

Abstract. C&C and deployment are base views for describing software architectures. Both views typically exist of several models that highlight particular aspects of the software. From our experience, managing the dependencies between architectural models is a challenging, yet crucial task for maintaining consistency of an architectural description. In this paper, we focus on the composition of architectural models and managing dependencies between models in such a composition. We formally define four basic relations between models and a composition operator that enables integration of several C&C models based on these relations. We elaborate on several semantic preserving properties of the composition operator, including traceability, consistency, and completeness. The relations and composition operator have been integrated in the xADL language and the AchStudio tool. We use excerpts of a distributed video-on-demand system designed with the tool as an illustrative case.

1 Introduction

The architecture of a software system defines its essential structures, which comprise software elements, the externally visible properties of those elements and the relationships between them [3]. A C&C view and deployment view are base views for any architectural description. Both views typically exist of several models that highlight particular aspects of the software. Since the models of in a view describe parts of the same system there are dependencies among the models. From our experience with designing non-trivial architectures, managing such dependencies remains a challenging, yet crucial task for maintaining consistency of an architectural description. During architectural design of several realistic case studies, including an industrial transportation system for logistic services [19][7] and a traffic control system [14], we encountered the following problems:

- Models often overlap, e.g. the same element appears in several models possibly in different roles; several elements overlap because a component in a model is refined in another model; or a number of models overlap with another model that describes an overview. Refraining from rigorously specifying such dependencies quickly leads to inconsistencies of the architectural description.
- Architects compose models to obtain a unified perspective on (parts of) the architecture, to understand the interactions between elements from different models, and to perform various types of analysis. However, stakeholders demand guarantees that a composed model is a correct representation of the models being composed. The lack of precisely defined dependencies among models makes it hard to compose models correctly and provide such guarantees.

- Tools support for describing dependencies between models is indispensably in practice. In particular, automating model composition would significantly increase the design comfort, allowing the architect to quickly obtain a unified perspective on the architecture and reveal conflicts between models. However, ambiguous descriptions of dependencies between models hampers the development of tool support.

In this paper, we put forward a formally founded approach to manage and exploit dependencies between C&C and deployment models. The formal specification paves the way to rigorous specification of dependencies between models and automation of model composition. Our particular focus is on the composition of C&C and deployment models and crucial properties of this composition. The approach is based on two key concepts: *relation* and *composition operator*. A relation defines a particular type of overlap between C&C models. We formally specify three common relations: unification, submodel, subelement and allocation. A composition operator defines the integration of several models based on the relations between these models, yielding a composed model. We define well-formedness rules for C&C models and relations, and we formally capture properties of model composition, namely traceability and semantic preservation. We use a formal language based on set theory that provides an accessible notation and compact specification of the core concepts. A complete formal description of the relations and the composition operator in Haskell, including proofs of the properties of model composition is available in [4].

In summary, the key contributions of this paper are: (1) a formal definition of three common *relations* and a *composition operator* for C&C models (section 2); (2) a formal specification of well-formedness rules (section 2) and a semantic specification of model composition (section 3). Section 5 discusses related work. Finally, section 6 discusses experiences and conclusions.

2 A formal model for relations and composition

Before going into the formal details, we introduce the main concepts used in this paper. Fig. 1 shows the concepts in the context of the conceptual model of ISO 42010. [1] The subset of concepts of ISO 42010 is shown in gray.

An *architectural description* (AD) is a collection of products to document a specific architecture. An AD is organized into one or more *architectural views*, where a view is a representation of a whole system from the perspective of a related set of concerns. A view is composed of one or more architectural models. A model contains a concrete description of architectural elements. This paper provides an in depth study of C&C views and deployment views. A *C&C view* is a base view to describe software systems and is an aggregation of one or more C&C models and relations. A *C&C model* contains concrete architectural elements like components, connectors and interfaces. A *deployment view* exists of infrastructure and deployment models. The infrastructure model describe the nodes, and connections between the nodes, on which the components can be deployed. A deployment model shows a concrete deployment of components. A *relation* describes a relation between two C&C models. Unification, submodel, subelement and allocation are specialization's of this relation. A *composition* integrates several C&C models using the relations among these models, defining an integrated model.

[1] ISO 42010 is based on ANSI/IEEE-Std-1471 and is currently undergoing standardization as the standard for "Systems and Software Engineering - Architectural Description".

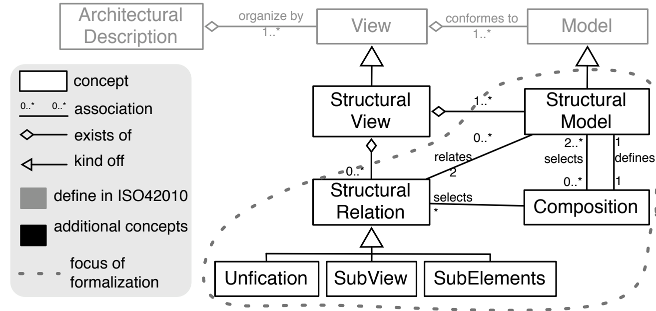

Fig. 1. Overview of concepts used in this paper.

The formalization is build up in steps: architectural models in 2.1, relations in 2.2 and composition in 2.3. In each step we briefly introduce the specific concepts, we explain the objective of the formalization, we formalize the concepts, and we define well-formedness rules.

2.1 Architectural models

In this section we define C&C, infrastructure and deployment models, their formalization model and well-formedness rules.

C&C model A *C&C model* consists of elements that have a runtime presence like components, connectors and interfaces, and links that describe the pathways of interactions. *Components* are the loci of computation in the architecture. A component is defined by a set of interfaces and an optional substructure. *Connectors* are the loci of communication in the architecture, specified in the same way as components. Connectors can only be used between components, but it is not mandatory to always define a connector between components. *Interfaces* are components' and connectors' portals to other components or connectors. The direction on an interfaces indicated the information flow, and can be in, out or inout. A *link* defines the topology of the architecture by connecting interfaces. A *substructure* defines the internal structures of either a component or a connector. A substructure consists of components, connectors and links. In the case of a substructure, *interface mappings* are used to map an outer interfaces (on the enclosing element) onto inner interfaces (on the internal element). Names of elements are fully qualified. For example, an interface `stream` in component `User` that is situated in model `VOD` has as name `VOD.User.stream`. When no confusion is possible, only the last part of the name is used to refer to the element, i.e. `stream`.

We use the architecture of a simplified Video on Demand (VOD) system as a running example. A VOD system allow users to search and watch videos at their demand. The extract used in this paper includes pay-per-view support, where a user pays for the videos she/he wants to see. The left hand side of Fig. 2 shows the VOD C&C model, defining the basic service to search and watch videos. The `User` component is responsible for interaction with the user. `Streamer` pushes streams to users based on commands it receives from `Query`. `Query` handles searches and requests from `User`, and`DB` manages the database of videos. Finally, the connector `SC` links `User` with `Streamer`, and `VQC` links `User` with `Query`.

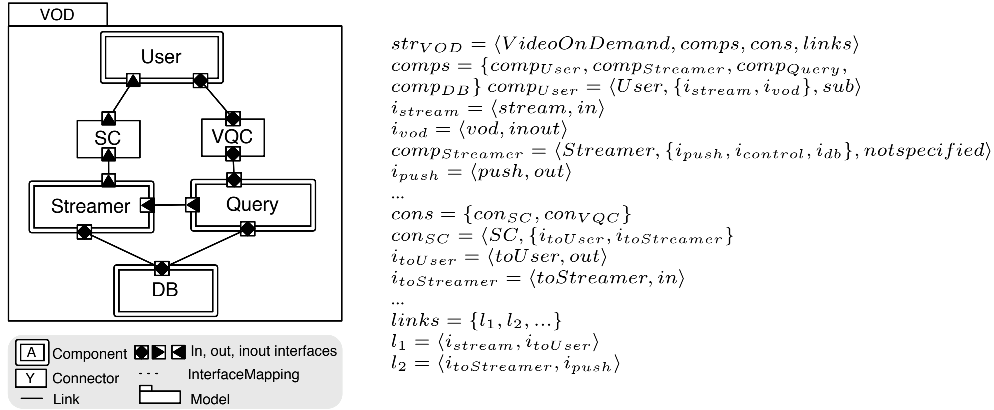

Fig. 2. Example C&C model. Left: visual notation. Right: extract of formal notation

Infrastructure model The *infrastructure model* is used to document the infrastructure available to the software system, before deciding on deployment (which happens in the deployment model). The model defines nodes and how they are connected to each other through communication channels. *Nodes* represent a computer system on which components can be deployed in the deployment model, and can be annotated with several properties to describe its characteristics. Example nodes for the TMS are embedded computer systems to control monitoring or signalization devices, application servers and databases servers. A *communication channel* represent the connections between the nodes. A communication channel exists of one or several *communication Links* that connect to the nodes itself. Example communication channels are a local area network (LAN), a wireless local area network (WLAN), a connection over the cell-phone network, etc.

Deployment model An *deployment model* is an extension of an infrastructural model to allow the allocation of components and connectors to nodes. A deployment model can be considered as a combined model[2], it combines both elements of the infrastructure and C&C models.

Formalization *Objective: providing a rigorous description of C&C models (1) to avoid ambiguity and inconsistency in descriptions of models; (2) to support the definition of relations and composition.*

The following sets define C&C models[3]:

[2] [9] states that a combined view is a view with a combination of styles. In the terminology of this paper this comes down to combining the types of several models.

[3] In this paper, we do not explicitly consider the difference between an element and a reference to an element. The extended version of the formalism [4] explicitly considers references.

ID	Set of names
$DIR = \{in, out, inout\}$	Set directions
$INT \subset ID \times DIR$	Set of interface
$CC \subset ID \times \mathbb{P}\, INT \times SUB$	Set of C&C
$COMP \subset CC$	Set of components
$CON \subset CC$	Set of connectors
$LINK \subset INT \times INT$	Set of Links
$IM \subset INT \times INT$	Set of interfacemapping
$SUB \subset \mathbb{P}\, COMP \times \mathbb{P}\, CON \times \mathbb{P}\, LINK \times \mathbb{P}\, IM$	Set of substructures
$none = \langle \emptyset, \emptyset, \emptyset, \emptyset \rangle \in SUB$	No substructure
$CCMODEL \subset ID \times \mathbb{P}\, COMP \times \mathbb{P}\, CON \times \mathbb{P}\, LINK$	Set C&C models

Since several well-formedness rules apply (see below), we use the subset construct to define several of the sets. We refer to an instance of an architectural element using its type name with its name as subscript, e.g. a component A is noted as $\texttt{comp}_A$. Furthermore, we use an abbreviate notation to refer to the internals of an element. For example, to refer to the interfaces of a component $\texttt{comp}_A$ we use the notation $\texttt{comp}_A\texttt{.ints}$. The following conventions on type names and internal parts are used:

$i_{name} = \langle dir \rangle$	Interface
$cc_{name} = \langle ints, sub \rangle$	CC
$comp_{name}, con_{name} = \langle ints, sub \rangle$	Component, connector
$l_i = \langle int1, int2 \rangle$	Link
$sub = \langle comps, cons, links, ims \rangle$	Substructure
$im_i = \langle inner, outer \rangle$	Interfacemapping
$ccmodel_{name} = \langle comps, cons, links \rangle$	Structural model

As an example, the left hand side of Fig. 2 shows an extract of the formalization of the C&C model shown on the right hand side. The formalization specifies the VOD model, the User and Streamer component with interfaces, the SC connector with interfaces, and two links that connect the two components through the connector.

The formalization of infrastructure and deployment models is similar:

$NODE \subset ID$	Set of nodes
$COMPATH \subset ID \times \mathbb{P}\, NODE$	Set of communication paths
$INFRA \subset ID \times \mathbb{P}\, NODE \times \mathbb{P}\, COMPATH$	Set of infrastructure models
$DNODE \subset NODE \times \mathbb{P}\, COMP \times \mathbb{P}\, CON \times \mathbb{P}\, LINK$	Deployment node
$DCOMPATH \subset ID \times \mathbb{P}\, NODE \times \mathbb{P}\, LINK$	Deployment compath
$DEPLOY \subset ID \times \mathbb{P}\, DNODE \times \mathbb{P}\, DCOMPATH$	Set of deployment models

This also involves naming conventions:

$node_{name}$	Node
$compath_{name} = \langle nodes \rangle$	Communication path
$infra_{name} = \langle nodes, paths \rangle$	Infrastructure model
$dnode_{name} = \langle comps, cons, links \rangle$	Node
$dcompath_{name} = \langle nodes, links \rangle$	Communication path
$deployment_{name} = \langle nodes, paths \rangle$	Infrastructure model

We define several functions. For the sake of simplicity, we define a function with the same name for several types (separated by a comma). This can be be translated to several functions for each of the types. [4] provides the full specification of these functions.

$ccs : CCMODEL, SUB \rightarrow \mathbb{P}\, CC$: returns all components and connectors in the given C&C model or substructure; we use the abbreviate notation $ccmodel.ccs$ and $sub.ccs$.

$allCC : CC, CCMODEL, \mathbb{P}CC, \mathbb{P}CCMODEL, DEPLOY, \mathbb{P}DEPLOY \rightarrow \mathbb{P}\, CC$: returns recursively all components and connectors in the given element (i.e. all elements in substructures are included).

$ints : CC, CCMODEL, \mathbb{P}CC, \mathbb{P}CCMODEL, DEPLOY, \mathbb{P}DEPLOY \rightarrow \mathbb{P}\, INT$: returns all interfaces that are in the given element (non recursively, so interfaces in substructures are not included).

$allInts : CC, CCMODEL, \mathbb{P}CC, \mathbb{P}CCMODEL, DEPLOY, \mathbb{P}DEPLOY \rightarrow \mathbb{P}\, INT$: returns recursively all interfaces in the given element.

$connected : INT \times INT \rightarrow Bool$: checks if two given interfaces are connected either by a link or by a link and several interface mappings.

Well-formedness rules The following rules exclude invalid and redundant tuples from the definition of C&C models:

$$\begin{gathered} \forall i_{name1}, i_{name2} \in INT : name1 = name2 \Rightarrow i_1 = i_2 \\ \forall c_{name1}, c_{name2} \in CC : name1 = name2 \Rightarrow c_1 = c_2 \\ \ldots \end{gathered} \tag{1}$$

$$\begin{gathered} \forall im \in IM : im.outer \neq im.inner \wedge im.outer.dir = im.inner.dir \\ \forall l \in LINK : (l.int1.dir = inout \wedge l.int2.dir = inout) \vee \\ (l.int1.dir = in \wedge l.int2.dir = out) \vee (l.int1.dir = out \wedge l.int2.dir = in) \end{gathered} \tag{2}$$

$$\begin{gathered} \forall ccm \in CCMODEL : \forall l \in ccm.links : \\ \{l.int1, l.int2\} \subseteq (ints\ ccm.comp \cup ints\ ccm.con) \\ \forall sub \in SUB, \forall l \in sub.links : \\ \{l.int1, l.int2\} \subseteq (ints\ sub.comp \cup ints\ sub.con) \end{gathered} \tag{3}$$

$$\begin{gathered} \forall sub \in SUB, \forall cc \in CC : cc.sub = sub \Rightarrow \\ (\forall im \in sub.im : im.inner \in (allInts\ sub.comp \cup allInts\ sub.con) \wedge \\ im.outer \in cc.ints) \end{gathered} \tag{4}$$

$$\begin{gathered} \forall ccm \in CCMODEL, \forall con \in ccm.con, \exists l \in struct.links : \\ (l.int1 \in con.ints \Rightarrow \exists comp \in ccm.comp : l.int2 \in comp.ints) \wedge \\ (l.int2 \in con.ints \Rightarrow \exists comp \in ccm.comp : l.int1 \in comp.ints) \end{gathered} \tag{5}$$

Rule 1 ensures that each element with a specific name is unique. The rule is only shown for interfaces and components. For connectors and C&C models, the definition is very similar. Rule 2 ensures that the directions of interfaces in interface mappings and links are compatible. Rule 3 and rule 4 make sure that links and interface mappings refer to interfaces of the correct C&C model or substructure. This rule excludes links between interfaces in different models and links between an interface in a substructure and an interface that is not in this substructure. Rule 5 makes sure that connectors can only be used between components. We only show the rule for C&C models, the definition for substructures is very similar.

The same uniqueness requirement as defined in rule 1 applies on all elements in the infrastructure and deployment models having a name.

$$\begin{gathered} \forall compath \in COMPATH : compath.nodes \neq \emptyset \\ \forall dcompath \in COMPATH : dcompath.nodes \neq \emptyset \end{gathered} \tag{6}$$

$$\begin{gathered} \forall infra \in INFRA, \forall compath \in infra.paths : \\ compath.nodes \subseteq infra.nodes \\ \forall deploy \in INFRA, \forall dcompath \in deploy.paths : \\ dcompath.nodes \subseteq deploy.nodes \end{gathered} \tag{7}$$

$$\begin{gathered} \forall dnode \in DNODE : \forall l \in dnode.links : \\ \{l.int1, l.int2\} \subseteq (ints\ dnode.comp \cup ints\ dnode.con) \end{gathered} \tag{8}$$

$$\begin{gathered} \forall deploy \in DEPLOY, \forall dcompath \in DCOMPATH, \forall l \in dcompath.links, \\ \exists dnode_1, dnode_2 \in deploy.nodes, \exists c_1 \in dnode_1.ccs, c_2 \in dnode_2.ccs : \\ dnode_1, dnode_2 \subseteq dcompath.nodes \wedge l.int1 \in c_1.ints \wedge l.int2 \in c_2.ints \end{gathered} \tag{9}$$

Rule 6 states that each communication path should connect nodes. Rule 7 states that the nodes in a communication path must be in the nodes of the respective infrastructure or deployment model. Rule 8 states that the links in a dnode must be to components who are also in this node. Rule 9 states the the links in a dcompath must be between components of the nodes connected by this path.

2.2 Relations

Relations document dependencies between different structural models. Explicitly specified relations allow an architect to see the implications of changing one model on another model. Relations form the basis for model composition what is subject of the next section. From our experience in several case studies, we deduced the following relations:

- **Unification**: expresses that two elements (either components or connectors) that appear in the two different C&C models are the same element. A unification is always companioned with interface unifications that express which interfaces are the same.
- **Submodel**: expresses that a C&C model describes the internal structure of an architectural element of another C&C model. Interface mappings are used to map internal elements on the enclosing element.
- **Subelements**: expresses that individual elements or groups of elements from one C&C model are part of the substructure of an element of another C&C model. Similar to submodel, the subelements relation uses interfaces mappings.
- **Allocation**: expresses that a set of component and connector is allocated onto a specific node from an infrastructure model.

Figure 3 illustrates the relations. The left hand side shows the C&C model `Accounting`. This model shows a connector that internally uses `Facade` to enforce that users pay to view a video. The right hand side shows the `VOD` model (as in Fig. 2). We visually illustrate the relations between elements with lines going from the one structure to the other structure. In the example, the bold dotted lines represent unifications, the light dotted lines unifications of interfaces. $\texttt{unif}_{User}$ expresses that `Accounting.User` and `VOD.User` are the same component. Even more, one of the interfaces of User is the same as expressed by $\texttt{iu}_{query}$. $\texttt{unif}_{VQC}$ expresses that `Accounting.AccountingCon` and `VOD.VQC` are the same connector, with two interface unifications $\texttt{iu}_{queryUser}$ and $\texttt{iu}_{queryServer}$.

Formalization.

Objective: providing a rigorous description of the three relations (1) to avoid ambiguity and inconsistency in documenting dependencies between C&C models; and (2) to support the definition of composition.

The following sets define the relations:

$$\begin{array}{rl} IU \subset INT \times INT & \text{Set of interface unifications} \\ UNIF \subset ID \times CC \times CC \times ID \times \mathbb{P}\, IU & \text{Set of unification relations} \\ SUBMODEL \subset ID \times CC \times CCMODEL \times \mathbb{P}\, IM & \text{Set of submodel relations} \\ ALLOC \subset ID \times NODE \times \mathbb{P}\, CC & \text{Set of allocations} \\ REL = UNIF \cup SUBMODEL \cup SUBELEM \cup ALLOC & \text{The set of relations} \end{array}$$

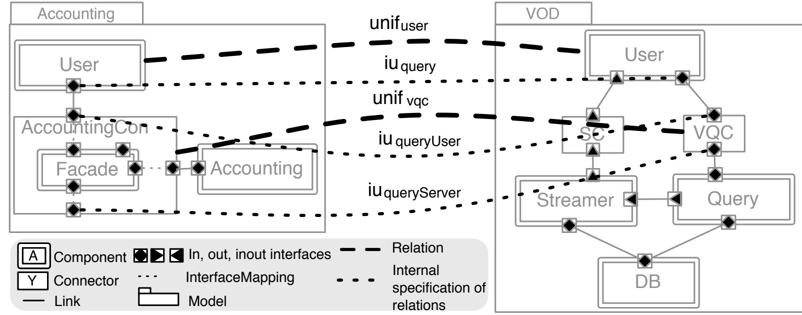

$$
\begin{aligned}
unif_{user} &= \langle comp_{VideoOnDemand.User}, comp_{Accounting.User}, User, ius_1 \rangle \\
ius_1 &= \{\langle i_{VideoOnDemand.User.query}, i_{Accounting.User.query} \rangle\} \\
unif_{vqc} &= \langle con_{Accounting}, con_{VQC}, VQC, ius_2 \rangle \\
ius_2 &= \{\langle i_{Accounting.service}, i_{VQC.query} \rangle, \langle i_{Accounting.server}, i_{Accounting.queryServer} \rangle\}
\end{aligned}
$$

Fig. 3. Example of unification relation. Top: two models with their relation. Bottom: corresponding formal specification of the relations.

$$
\begin{aligned}
iu &= \langle ui1, ui2, newName \rangle && \text{Interface unification} \\
unif_{name} &= \langle elem1, elem2, newName, ius \rangle && \text{Unification} \\
subelem_{name} &= \langle target, subelems, ims \rangle && \text{Subelem} \\
submodel_{name} &= \langle target, submodel, ims \rangle && \text{Submodel} \\
alloc_{name} &= \langle node, ccs \rangle && \text{Allocation}
\end{aligned}
$$

Fig. 3 contains a formal description of the relations in our running example.

Well-formedness rules. The same uniqueness requirement as defined in rule 1 applies for the relations. Furthermore, the function `allCC` is also defined for relations.

$$
\begin{gathered}
\forall unif \in UNIF, \exists s_1, s_2 \in CCMODEL : \\
(unif.elem1 \in s_1.comps \wedge unif.elem2 \in s_2.comps) \vee \\
(unif.elem1 \in s_1.cons \wedge unif.elem2 \in s_2.cons)
\end{gathered} \tag{10}
$$

$$
\begin{gathered}
\forall unif \in UNIF, \forall s_1, s_2 \in CCMODEL : \\
unif.elem1 \in s_1.ccs \wedge unif.elem2 \in s_2.ccs \Rightarrow s_1 \neq s_2
\end{gathered} \tag{11}
$$

$$
\begin{gathered}
\forall unif \in UNIF, \forall iu \in unif.ius : \\
(iu.ui1 \in unif.elem1.ints \wedge iu.ui2 \in unif.elem2.ints)
\end{gathered} \tag{12}
$$

$$
\begin{gathered}
\forall sub \in SUBELEM, \forall cc \in sub.subelems, \exists s_1, s_2 \in CCMODEL : \\
sub.target \in s_1.ccs \wedge cc \in s_2.ccs \wedge s_1 \neq s_2
\end{gathered} \tag{13}
$$

$$
\begin{gathered}
\forall submodel \in SUBMODEL : \\
submodel.target \notin (allCC\ submodel.submodel)
\end{gathered} \tag{14}
$$

$$
\begin{gathered}
\forall submodel \in SUBMODEL, \forall im \in submodel.ims, \\
\exists cc \in submodel.submodel.ccs : \\
im.outer \in submodel.target.ints \wedge im.inner \in cc.ints
\end{gathered} \tag{15}
$$

Rule 10 ensures that unifications between a component and a connector are not possible. Rule 11 ensures that two unified elements are always from a two different C&C models. Rule 12 ensures that all interface unifications are between the interfaces of the unified elements. Rule 13 ensures the target element and the subelements in a subelements relation are from different models. Rule 14 ensures that the target element is not an element of the submodel. Rule 15 ensures that interface mappings in relations are indeed between the related elements. There are no specific well-formedness rules for the internals of allocation links.

2.3 Composition

Composition defines the integration of several models based on the relations defined between these models. We distinguish between the *composition specification* and the *composition operator*. The composition specification *selects* models and relations. The composition operator takes a specification and *defines* an integrated model and a set of *traces*. The set of traces relate all elements in the output model of a composition operator with the elements in the input models selected by the composition specification.

Formalization.

Objective: providing a rigorous description of model composition and semantic preserving properties of the composition (1) to provide guarantees to the stakeholders that a composed model is a correct representation of the input models, (2) to support various types of analysis, and (3) to pave the way to automate model composition.

We start with a formal model of the composition specification:

$$COMPSPEC \subset ID \times \mathbb{P}\, MODEL \times \mathbb{P}\, REL \times ID \qquad \text{Set of composition specs}$$
$$spec_{name} = \langle inModel, inRels, modelName \rangle \qquad \text{Naming conventions}$$

We use the terms 'input models' (`inModel`) and 'input relations' (`inRels`) to refer to the models and relations selected in the composition specification. Fig. 4 shows an example of a composition specification based on the models and unification relations defined in Fig. 3.

Several well-formedness rules apply:

$$\begin{gathered}\forall spec \in COMPSPEC, \forall rel_1, rel_2 \in spec.inRels : \\ rel_1, rel_2 \in SUBMODEL \cup UNIFICATION \wedge \\ rel_1 \neq rel_2 \Rightarrow (allCC\; rel_1 \cap allCC\; rel_2 = \emptyset)\end{gathered} \tag{16}$$

$$\begin{gathered}\forall spec \in COMPSPEC, \forall rel_1, rel_2 \in spec.inRels : \\ rel_1, rel_2 \in ALLOC \wedge \\ rel_1 \neq rel_2 \Rightarrow (allCC\; rel_1 \cap allCC\; rel_2 = \emptyset)\end{gathered} \tag{17}$$

$$\begin{gathered}\forall spec \in COMPSPEC, \forall rel \in spec.inputRels : \\ allCC\; rel \subseteq allCC\; spec.inputModels \wedge \\ allInts\; rel \subseteq allInts\; spec.inputModels\end{gathered} \tag{18}$$

Rule 16 ensures there is no overlap between two relations between C&C models. Currently, we require that there is no overlap between the relations. This could be relaxed in the future, but we need to do further investigate the interplay between the relations. Rule 17 ensures there is no overlap between allocation relations. Rule 18 ensures that all relations are within the scope of the input models.

Next, we define a composition operator. The composition operator takes a composition specification and defines an integrated model and a set of traces. This leads to a function with the following signature:

$$compositionFunction : COMPSPEC \rightarrow MODEL \times \mathbb{P}\, TRACE$$

We use the term 'output model' (`outStr`) to refer to the model defined by the composition; `traces` denotes the resulting set of traces. Note that the composition function already implies that the output model is well-formed since it is an element of STRUCT for which the well-formedness rules apply. Each trace contains two elements, an element in the output model and an element in the input model. Traces allow us to keep an overview of the interconnections between the input models and the output model. Traces are defined as follows:

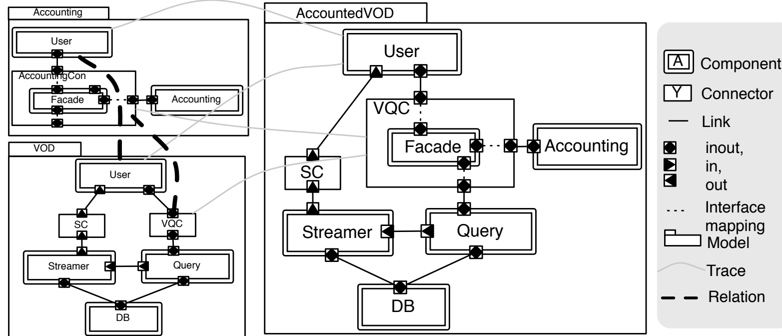

$$spec_{AccountedVOD} = \langle AccountedVOD, \{str_{VOD}, str_{Accounting}\}, \{unif_{user}, unif_{vqc}\}\rangle$$

Fig. 4. Example of composition. Left: the two input models, related by two unifications. Right: the output model. Gray lines from left to right show example traces for the unified elements.

$$TRACES = (MODEL \times MODEL) \cup (CC \times CC) \cup (INT \times INT)$$
$$trace = \langle out, in \rangle$$

The operational description of the composition operator is too complex to show in this paper; [4] provides all the details. Here, we briefly explain the three main steps of model composition and illustrate the result using the running example. In section 3, we define the semantics of the composition operator by formally capturing information preserving properties.

The left hand side of Fig. 4 shows the two input models (`Accounting` and `VOD`), the right hand side shows the resulting output model (`AccountedVOD`). The thin gray lines from left to right represent the traces for components and connectors (to simplify the figure, we left out the traces for interfaces). The main steps of model composition are:

1. Apply the relations one by one. The order is not relevant since relations are independent (implied by rule 16 and rule 17). Applying means: (i) use the information in the relation to define an integrated architectural element and add this element in the output model, and (ii) add traces between the respective elements in the input models and the output model. For example, applying the relation $\texttt{unif}_{User}$ yields: (i) the `User` component with two interfaces in the output model, and (ii) two traces for the `User` component from the output model to two `User` components in the input models, and a set of traces for the interfaces. Similarly, relation $\texttt{unif}_{VQC}$ adds the `VQC` connector and its interfaces, and a set of traces for all elements. Notice that the internal structure of `Accounting.AccountingCon` is preserved by the unification.
2. For each component and connector that can not be found in the traces: copy the element from the input models to the output model, and add the respective traces for the element. For our running example, this step adds the elements `Accounting`, `Streamer`, `Query`, `DB`, `SC` to the output model, and traces for all the elements and their internals.
3. Link the components and connectors in the output model. Each link in the input models is translated to a link in the output model using the traces.

3 Semantic-preserving properties of the composition operator

The formal specification of model composition allows us to formally capture semantic preservation properties of the composition operator. Semantic-preserving properties provide guarantees to the stakeholders that the output model is a correct representation of the composed models. We discuss four properties: traceability, consistency, model completeness, and relational completeness. For each property we give a brief introduction, we explain the objective of the formalization, and we formally define the property. Proofs of the properties can be found in [4].

3.1 Traceability

Traceability has widely been recognized as an important aspect of software development. IEEE [15] defines traceability as "the degree to which a relationship can be established between two or more products of the development process, especially products having a predecessor-successor or master-subordinate relationship; for example, the degree to which the requirements and design of a given software component match." The composition operator presented in this paper supports traceability between the input models and the output model of a composition.

Formalization.

Objective: ensuring that all elements from the output model are traceable to elements from the input models (1) to enable architects to trace back the origin of each element of a composed model; (2) to provide the basis for describing and proving the other semantic preserving properties, and (3) to pave the way for automatic analysis, such as consistency checking and impact analysis.

The formal definition of traceability is captured as follows:

$$\forall model_i \in inModels : \langle outModel, model_i \rangle \in traces \tag{19}$$

$$\begin{gathered}\forall c_{out} \in allCC\ outModel, \exists model \in inModels, \exists c_{in} \in allCC\ model : \\ \langle c_{out}, c_{in} \rangle \in traces\end{gathered} \tag{20}$$

$$\begin{gathered}\forall node_{out} \in outModel.nodes, \exists model \in inModels, \exists node_{in} \in model.nodes : \\ \langle node_{out}, node_{in} \rangle \in traces\end{gathered} \tag{21}$$

$$\begin{gathered}\forall compath_{out} \in outModel.path, \exists model \in inModels, \exists compath_{in} \in model.path : \\ \langle compath_{out}, compath_{in} \rangle \in traces\end{gathered} \tag{22}$$

$$\begin{gathered}\forall c_{out} \in (allCC\ outModel), \forall i_{out} \in c_{out}.ints, \exists model \in inModels, \\ \exists c_{in} \in (allCC\ model), \exists i_{in} \in c_{in}.ints : \\ \langle c_{out}, c_{in} \rangle \in traces \wedge \langle i_{out}, i_{in} \rangle \in traces\end{gathered} \tag{23}$$

$$\begin{gathered}\forall i_{out1}, i_{out2} \in (allInts\ outModel), \forall model \in inModels, \\ \exists i_{in1}, i_{in2} \in (allInts\ model) : (\langle i_{out1}, i_{in1} \rangle \in traces \wedge \\ \langle i_{out2}, i_{in2} \rangle \in traces\ \wedge connected\ i_{out1}\ i_{out2}) \Rightarrow connected\ i_{in1}\ i_{in2}\end{gathered} \tag{24}$$

Equation 19 states that there is a trace between the output model and each input model of a model composition. Equation 20 states that for each component and connector in the output model, there exists a corresponding component or connector in the input models, identifiable through traces. Equation 21 and equation 22 state the same for nodes and compaths respectively. Equation 23 states that for each interface of each component and connector in

the output model, there exists an interface of the corresponding component or connector in the input models, identifiable through traces. One could interpret traces as a 'binary surjective relation' between the set of all elements with a name in the output model, and the set of all elements with a name in the input models. Equation 24 states that each connection in the output model can be traced back to a connection in the input model. Two interfaces are connected if there is a link between the interfaces, or between interface mappings of the interfaces (captured in the connected function defined in section **??**). Note that links are not directly traceable. Traces of links are implied by the connections.

3.2 Consistency

Consistency refers to the compatibility of the output model with the input models. In particular, the output model of a model composition should not contradict the input models. We define two consistency properties: (1) two separate components of an input model cannot become a single component in the output model, and (2) if an input model defines that an element is part of a substructure, it must be part of the same substructure in the output model.

Formalization.

Objective: ensuring that the output model is consistent with the input models (1) to make the guarantees of consistency of model composition explicit; (2) to enable automation of consistency checks.

Model composition must preserve the distinction between elements defined in the input models. For example, if an input model contains components A and B, it is not allowed that the output model contains a single component that represents both A and B. This would imply that the composition does not respect the distinction between components A and B in the input model. The same holds for interfaces. This is captured in the following two formal statements:

$$\begin{gathered}\forall c_{out} \in allCC\ outModel, \forall model_1, model_2 \in inModels,\\ \exists c_1 \in allCC\ model_1, \exists c_2 \in allCC\ model_2 :\\ (\langle c_{out}, c_1\rangle \in traces \wedge \langle c_{out}, c_2\rangle \in traces \wedge c_1 \neq c_2) \Rightarrow model_1 \neq model_2\end{gathered} \tag{25}$$

$$\begin{gathered}\forall node_{out} \in outModel.nodes, \forall model_1, model_2 \in inModels,\\ \exists node_1 \in model_1.nodes, \exists node_2 \in model_2.nodes :\\ (\langle node_{out}, node_1\rangle \in traces \wedge \langle node_{out}, node_2\rangle \in traces \wedge node_1 \neq node_2)\\ \Downarrow\\ model_1 \neq model_2\end{gathered} \tag{26}$$

$$\begin{gathered}\forall compath_{out} \in outModel.path, \forall model_1, model_2 \in inModels,\\ \exists compath_1 \in model_1.path, \exists compath_2 \in model_2.path :\\ (\langle compath_{out}, compath_1\rangle \in traces \wedge\\ \langle compath_{out}, compath_2\rangle \in traces \wedge node_1 \neq node_2)\\ \Downarrow\\ model_1 \neq model_2\end{gathered} \tag{27}$$

$$\begin{gathered}\forall c_{out} \in allCC\ outModel, \forall i_{out} \in c_{out}.ints, \exists model_1, model_2 \in inModels,\\ \exists c_1 \in (allCC\ model_1), \exists c_2 \in (allCC\ model_2),\\ \exists i_1 \in c_1.ints, \exists i_2 \in c_2.ints : \langle c_{out}, c_1\rangle \in traces \wedge \langle c_{out}, c_2\rangle \in traces \wedge\\ \langle i_{out}, i_1\rangle \in traces \wedge \langle i_{out}, i_2\rangle \in traces \wedge i_1 \neq i_2\\ \Downarrow\\ model_1 \neq model_2\end{gathered} \tag{28}$$

Equation 25 states that for each component and connector in an output model, if there are two corresponding components or connectors in the input models, these components or connectors must be from different input models. Equation 26 and equation 27 state the same for nodes and compaths respectively. Equation 28 states that for each interface of a component or connector in the output model, if there are two corresponding interfaces in the input models, these interfaces must be from different input models.

Next to preservation of the differences between elements, substructures must be preserved. This is captured in the following equations:

$$\forall c_{in} \in inStrs.ccs, \forall c_{insub} \in c_{in}, \nexists c_{out} \in outStr.ccs : \langle c_{out}, c_{in} \rangle \in traces \tag{29}$$

$$\begin{gathered} \forall c_{in} \in inStrs.ccs, \forall c_{insub} \in c_{in}, \exists c_{out} \in outStr.ccs, \exists c_{outsub} \in c_{out}.sub.ccs : \\ \langle c_{out}, c_{in} \rangle \in traces) \Rightarrow \langle c_{out}, c_{in} \rangle \in traces \end{gathered} \tag{30}$$

Equation 29 states that if an element is part of a substructure in the input models, there exists no corresponding element that is not part of a substructure in the output model. Equation 30 states that if an element is part of a substructure of an input element c_{in}, and there exists a corresponding element part of a substructure of output element c_{out}, there must be a trace between c_{in} and c_{out}. This makes sure that the element is part of the same substructure in the output model.

3.3 Model completeness

Model completeness refers to the inclusion of all information of the input models in the output model of a model composition.

Formalization.

Objective: ensuring that the output model of a model composition covers all information covered in the input models (1) to make the guarantees of semantic preserving of input models explicit; (2) to enable automation of completeness checks.

Model completeness requires that for each element in the input models there must be a corresponding element in the output model, identifiable through traces. Together with the consistency property (section 3.2) this implies that the output model covers all information of the input models. Model completeness is captured in the following formal statements:

$$\begin{gathered} \forall model_{in} \in inModels, \forall c_{in} \in allCC\ model_{in}, \\ \exists! c_{out} \in allCC\ outModel : \langle c_{out}, c_{in} \rangle \in traces \end{gathered} \tag{31}$$

$$\begin{gathered} \forall model_{in} \in inModels, \forall node_{in} \in model_{in}.nodes, \\ \exists! node_{out} \in outModel.node : \langle node_{out}, node_{in} \rangle \in traces \end{gathered} \tag{32}$$

$$\begin{gathered} \forall model_{in} \in inModels, \forall compath_{in} \in model_{in}.path, \\ \exists! compath_{out} \in outModel.path : \langle compath_{out}, compath_{in} \rangle \in traces \end{gathered} \tag{33}$$

$$\begin{gathered} \forall model \in inModels, \forall c_{in} \in allCC\ model, \forall i_{in} \in ccInts\ c_{in}, \\ \exists! c_{out} \in allCC\ outModel, \exists! i_{out} \in ccInts\ c_{out} : \\ \langle c_{out}, c_{in} \rangle \in traces \wedge \langle i_{out}, i_{in} \rangle \in traces \end{gathered} \tag{34}$$

$$\begin{gathered} \forall model \in inModels, \forall c_{in} \in (allCC\ model), \exists! c_{out} \in (allCC\ outModel) : \\ \langle c_{out}, c_{in} \rangle \in traces \wedge (c_{in}.sub = notspecified \vee \\ (\forall c_{insub} \in allCC\ c_{in}, \exists! c_{outsub} \in allCC\ c_{out} : \\ \langle c_{outsub}, c_{insub} \rangle \in traces)) \end{gathered} \tag{35}$$

$$\begin{gathered}\forall model \in inModels, \forall i_{in1}, i_{in2} \in (allInts\; model),\\ \exists! i_{out1}, i_{out2} \in (allInts\; outModel):\\ \langle i_{out1}, i_{in1}\rangle \in traces \wedge \langle i_{out2}, i_{in2}\rangle \in traces \wedge connected\; i_{in1}\; i_{in2}\\ \Downarrow\\ connected\; i_{out1}\; i_{out2}\end{gathered} \tag{36}$$

Equation 31 states that for each component and connector in the input models, there exist a unique counterpart in the output model, identifiable through traces. Equation 32 and equation 33 state the same for nodes and communication paths. Equation 34 states that for each interface of each component and connector in the input models, there exists a corresponding interface in the uniquely corresponding component or connector of the output model. Equation 35 states that for each component and connector in the input models, the corresponding component or connectors in the output model has a corresponding substructure. Equation 36 states that for each connection between interfaces in the input models, there exists a connection between the corresponding interfaces in the output model. Notice that the completeness property implies that the set of traces can be interpreted as a total surjective function between all elements in the input models with a name and all elements in the output model with a name.

3.4 Relational completeness

Relational completeness refers to the inclusion of all information of the relations between the input models in the output model of a model composition.

Formalization.

Objective: ensuring that the output model of a model composition covers all information covered in the relations between the input models (1) to make the guarantees of semantic preserving of relations between input models explicit; (2) to enable automation of completeness checks.

The following equations formalize relational completeness. The equations have to be considered in the context of the already defined properties:

$$\begin{gathered}\forall unif_i \in inRel \cap UNIF, \exists! cc_{out} \in (allCC\; outStr), \exists t_1, t_2 \in traces:\\ \begin{pmatrix} t_1 = \langle cc_{out}, unif_i.elem1\rangle \wedge t_2 = \langle cc_{out}, unif_i.elem2\rangle \wedge\\ (\nexists t_3 \in trace : t_3.out = cc_{out} \wedge t_3 \neq t_1 \wedge t_3 \neq t_2)\end{pmatrix} \wedge\\ \begin{pmatrix} \forall iu_j \in unif_i.ius, \exists! i_{out} \in cc_{out}.ints, \exists t_1, t_2 \in traces:\\ t_1 = \langle i_{out}, iu_j.ui1\rangle \wedge t_2 = \langle i_{out}, iu_j.ui2\rangle \wedge\\ (\nexists t_3 \in trace : t_3.out = i_{out} \wedge t_3 \neq t_1 \wedge t_3 \neq t_2)\end{pmatrix} \wedge\\ \begin{pmatrix} \forall i_{in} \in (unif_i.elem1.ints \cup unif_i.elem2.ints)/allInts\; unif_i.ius,\\ \exists i_k \in cc_{out}.ints : \langle i_k, i_{in}\rangle \in traces\end{pmatrix}\end{gathered} \tag{37}$$

$$\begin{gathered}\forall submodel_i \in inRel \cap SUBMODEL, \exists! cc_{out} \in (allCC\; outStr):\\ \langle cc_{out}, submodel_i.target\rangle \in traces \wedge\\ (\forall cc \in submodel_i.submodel.ccs, \exists! cc_{subout} \in cc_{out}.sub : \langle cc_{subout}, cc\rangle \in traces)\end{gathered} \tag{38}$$

$$\begin{gathered}\forall subelem_i \in inRel \cap SUBELEM, \exists! cc_{out} \in (allCC\; outStr):\\ \langle cc_{out}, subelem_i.target\rangle \in traces \wedge\\ (\forall cc \in subelem_i.subelems, \exists! cc_{subout} \in cc_{out}.sub : \langle cc_{subout}, cc\rangle \in traces)\end{gathered} \tag{39}$$

$$\begin{gathered}\forall str \in inStrs, \forall cc_{in} \in str.ccs/allCC\; inRel, \exists! cc_{out} \in strOut, \exists t_{c1} \in traces:\\ t_{c1} = \langle cc_{out}, cc_{in}\rangle \wedge (\nexists t_{c2} \in traces : t_{c2}.out = cc_{out} \wedge t_{c2} \neq t_{c1}) \wedge\\ \begin{pmatrix} \forall i_{in} \in cc_{in}.ints, \exists! i_{out} \in cc_{out.ints}, t_{i1} \in traces : t_{i1} = \langle i_{in}, i_{out}\rangle \wedge\\ (\nexists t_{i2} \in traces : t_{i2}.out = i_{out} \wedge t_{i2} \neq t_{i1})\end{pmatrix}\end{gathered} \tag{40}$$

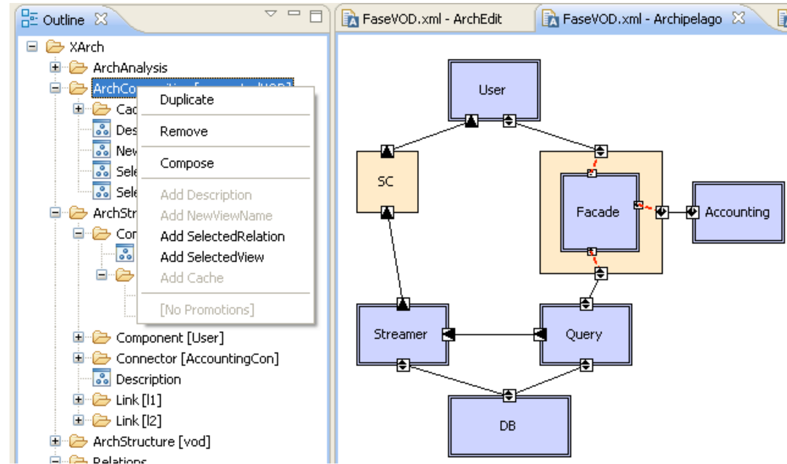

Fig. 5. Snapshot of the tool when the composition of AccountedVOD is triggered (as specified in Fig. 4)

Equation 37 defines what information is preserved from a unification. The first part of the equation (lines 2-3) captures the semantic preservation with respect to components and connectors. The second part (lines 4-6) captures semantic preservation of unified interfaces. The last part (lines 7-8) captures semantic preservation for non-unified interfaces. Equation 38 defines what information is preserved for a submodel relation; equation 39 defines semantic preservation for a subelems relation. Finally, equation 40 states that all components and connectors (and their interfaces) that are not part of a relation in the input moels must be uniquely traceable to an element in one of the output model. This last equation partly overlaps with the definition of completeness, but with the additional requirement that there is a *unique* trace.

4 xADLComposition: a tool for relations and composition in xADL

To make model composition practical, tools support is indispensably. Tool support allows architects to quickly obtain a unified perspective on the architecture and reveal conflicts between models. An unambiguous description of relations and composition is a prerequisite for developing proper tool support. The formalization presented in the previous sections provides such unambiguous description and served as a necessary basis for the development of the tools. Moreover, the semantic preserving properties explicitly describe guarantees about the composition operator.

We integrated the relations and the composition operator in ArchStudio [5], providing a Java implementation of the relations and composition in the xADL language [10]. The extension manifests itself in two aspects: a *language extension* introducing relations and composition, and an extension of the tool to enable *automatic composition*. We also extended the *visual tool* to show substructures. In the current version of the tool, composition is triggered manually generating the resulting output model and the accompanying set of traces. Fig. 5 shows a snapshot of the tool.

The formal specification of relations and composition provided a sound basis to guide the designers and developers building a consistent and conceptually sound tool. Especially the introduction of traces proved to be a very useful. During composition, the set of traces

always contains an up-to-date mapping between the output model and input models of all elements that are already processed by the composition operator. This allows to go back and forth between output and input and check whether elements are already processed during composition.

Several extensions of the tool are planned as future work. (1) We plan to add a visual and intuitive interface to specify relations. (2) Currently, the tool requires explicit triggering to start a composition. We plan to enable automatical synchronization of a composed model if one of the input models changes. (3) Finally, we plan to include support for automatic consistency checks. In particular, we plan to develop tool support to automatically check compatibility of relations between models. More information about the tool can be found on [5]. The website includes several examples of composition, screenshots and demonstration videos, and the source code of the tool.

5 Related work

In previous work we performed a broad study of relations between views and models and proposed a conceptual framework to characterize the relations in terms of usage, scope and mechanism [6]. Here we focus on a number of closely related approaches for composition.

Abi-Antoun et al. [1] describe an algorithm and tool for automatic differencing and merging models, focussing on reconciling different version of the same model. Contrary to our approach relations are not manually specified, but merging is done based on names and structural similarities which are very common for different versions of the same model. Automatic identification of relations based on names could be of interest for our approach too. Architectural stratification [2] relates several architectural strata (C&C models in ISO 42010 terminology). Each stratum represents the architectural structure on a certain level of abstraction, and step-wise refinement is used to specify the strata. In each step, a refinement transformation is applied that breaks up a connector in several elements introducing a particular concern. The submodel and subelements relations in our approach are inspired by this work. The approach presented in this paper is not limited to refinement and allows automated integration of models based on the relations between the models. Several authors focus on the integration of several types of models. In [12, 11], Egyed et al. put forward a view integration framework to enable automated design analysis of several heterogeneous view (models in ISO 42010 terminology). The main contribution is on transforming the information from heterogeneous views to allow a comparison, for example to check consistency. Ehrig et al. [13] use views in visual languages and graph morphisms to integrate different views, focussing on extending graph morphisms to easily allow introducing new model types. The focus of the research in this paper is not on heterogeneous model types, but on an in depth study of composition of C&C models, including a formal proof of the properties of composition and tool support [4].

As a final remark, there is a relation between the work described in this paper and aspect orientation and model transformation. For aspect orientation, composition (often referred to as weaving) is used to integrate crosscutting concerns with the basis system. Some aspect oriented techniques for architecture (e.g. [16]), detailed design (e.g. Theme/UML [8]), and programming (e.g. HyperJ [18]) use similar ways of composing structures. For model transformation, OMG recently defined a a standard called the Query/View/Transformation Specification (QVT) [17]. QVT advocates specifying relations explicitly and use the relations in transformation definitions. Our approach can be considered as one concrete instantiation of the QVT model, where a transformation corresponds to a composition of C&C models.

6 Discussion and conclusion

Managing the dependencies between architectural models remains a challenging yet crucial task for maintaining consistency of an architectural description. In this paper we have put forward a formally founded approach to manage and exploit such dependencies. We formalized *C&C models*, three common *relations* and a *composition operator* for C&C models. We defined well-formedness rules for C&C models and relations, and we formally captured semantic preserving properties of model composition, namely traceability, consistency and completeness. As a proof of concept, we have built a demonstration tool in ArchStudio that supports the relations and automated composition in the xADL language.

This paper achieves the following : (1) the formal description of C&C models avoids ambiguity and inconsistency of descriptions of *C&C models*; (2) the formal description of the three relations avoids ambiguity and inconsistency in documenting *dependencies between* C&C models; and (3) the formal description of model composition and semantic preserving properties of the composition make *explicit* the *guarantees* of the composition operator. Together, the formal description of (1-3) paved the way to *automate composition* in a tool.

The work in this paper fits in a larger project to support relations and composition of multiple views. In this paper, we have focussed on an in depth study of C&C and deployment views. From our experience, formalization and providing tool support for relations and composition, and in particular providing guarantees about model composition, is non-trivial. But, obviously relations also exist between different views and different types of models. Currently, we are broadening our research to different types of models. We already performed first steps towards supporting a deployment relation between a component & connector model and a runtime infrastructure model to generate a deployment view. Initial experiments are promising. Support for relations with statecharts is currently also under development.

References

1. M. Abi-Antoun, J. Aldrich, N. Nahas, B. Schmerl, and D. Garlan. Differencing and merging of architectural views. *Automated Software Engineering.*, 15(1):35–74, 2008.
2. C. Atkinson and T. Kühne. Aspect-oriented development with stratified frameworks. *IEEE Software*, 20(1):81–89, 2003.
3. L. Bass, P. Clements, and R. Kazman. *Software Architectures in Practice*. Addison, 2003.
4. N. Boucké. Formal proofs of well-formedness and information-preservation properties and haskell source code. `http://www.cs.kuleuven.be/~nelis/composition.htm`.
5. N. Boucké. xADLComposition: a tool for view composition in xADL. `http://www.cs.kuleuven.be/~nelis/xADLComposition`.
6. N. Boucké, D. Weyns, R. Hilliard, and T. Holvoet. Characterizing relations between architectural views. *Proceedings European Conference on Software Architecture*, LNCS 5292:66–81, 2008.
7. N. Boucké, D. Weyns, K. Schelfthout, and T. Holvoet. Applying the ATAM to an architecture for decentralized control of a transportation system. In *Quality of Software Architectures conference (QoSA)*, volume LNCS 4214, 2006.
8. S. Clarke and E. Baniassad. *Aspect-Oriented Analysis and Design*. Addison-Wesley, 2005.
9. P. Clements, F. Bachman, L. Bass, D. Garlan, J. Ivers, R. Little, R. Nord, and J. Stafford. *Documenting Software Architectures, Views and Beyond*. Addison Wesley, 2003.
10. E. Dashofy, A. van der Hoek, and R. Taylor. A comprehensive approach for the development of modular software architecture description languages. *ACM Transactions on Software Engineering and Methodology (TOSEM)*, 14(2):199–245, 2005.
11. A. Egyed. *Heterogeneous view integration and its automation*. PhD thesis, Los Angeles, CA, USA, 2000. Adviser-Barry William Boehm.

12. A. Egyed and N. Medvidovic. A formal approach to heterogeneous software modeling. In *FASE*, pages 178–192, 2000.
13. H. Ehrig, K. Ehrig, C. Ermel, and U. Prange. Consistent integration of models based on views of visual languages. In *Fundamental Approaches to Software Engineering*. LNCS, 2008.
14. A. Helleboogh, D. Weyns, T. Holvoet, and N. Boucké. On adls and tool support for documenting view-based architectural descriptions. In *Fourth SEI Software Architecture Technlogy User Network Workshop, SATURN*, 2008.
15. IEEE. IEEE standard glossary of software engineering terminology. *IEEE Std 610.12*, Dec 1990.
16. M. Kandé. *A Concern-Oriented Approach to Software Architecture*. PhD thesis, École Polytechnique Fédérale de Lausanne, 2003.
17. OMG. *Meta Object Facility 2.0: Query/View/Transformation Specification*, August 2007.
18. P. Tarr, H. Ossher, W. Harrison, and S. Sutton. N degrees of separation: Multi-dimensional separation of concerns. In *Int. Conf. on Software Engineering*, pages 107–119, 1999.
19. D. Weyns, T. Holvoet, K. Schelfthout, and J. Wielemans. Decentralized control of automatic guided vehicles: Applying multi-agent systems in practice. In *Development Track OOPLSA, Int. Conf. on Object-Oriented Programming, Systems, Languages, and Applications*, 2008.