

# A service-oriented UML profile with formal support <sup>★</sup>

Roberto Bruni<sup>1</sup>, Matthias Hölzl<sup>3</sup>, Nora Koch<sup>2,3</sup>, Alberto Lluch Lafuente<sup>1</sup>,  
Philip Mayer<sup>3</sup>, Ugo Montanari<sup>1</sup>, and Andreas Schroeder<sup>3</sup>

<sup>1</sup>University of Pisa, Italy

<sup>2</sup>Cirquent GmbH, Germany

<sup>3</sup>Ludwig-Maximilians-Universität München, Germany

**Abstract.** We present a UML Profile for the description of service oriented applications. The profile focuses on style-based design and reconfiguration aspects at the architectural level. Moreover, it has a formal support in terms of an approach called Architectural Design Rewriting, which enables formal analysis of the UML specifications. We show how our prototypical implementation can be used to analyse and verify properties of a service oriented application.

## 1 Introduction

Service-oriented computing is a paradigm centered around the notion of service: autonomous, platform-independent computational entities that can be described, published, discovered, and dynamically assembled for developing massively distributed, interoperable, evolvable systems and applications. However, services are still developed in a poorly systematic, ad-hoc way. Full fledged theoretical foundations are missing, but they are urgently needed for achieving trusted interoperability, predictable compositionality, and for guaranteeing security, correctness, and appropriate resource usage.

The IST-FET Integrated Project *SENSORIA* aims at developing a comprehensive approach to the engineering of service-oriented software systems where foundational theories, techniques and methods are fully integrated into pragmatic software engineering processes. The development of mathematical foundations and mathematically well-founded engineering techniques for service-oriented computing constitutes the main research activity of *SENSORIA*.

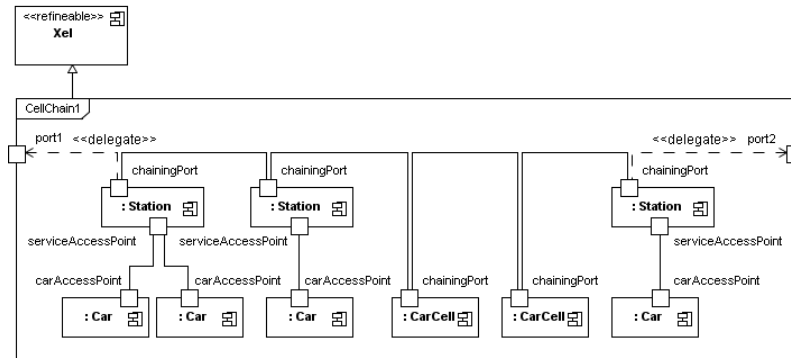
In this paper we present recent efforts within *SENSORIA* aimed to develop high-level modelling languages with strong formal support. More precisely, we present a novel extension of UML4SOA, our UML2 profile and define here its formal semantics. The presentation is illustrated on a simple example taken from the automotive domain.

UML4SOA is an extension of the UML2, the *lingua franca* of object-oriented software analysis and design. UML4SOA enhances UML2 with concepts for modelling structural and behavioural aspects of services. In this paper we present for the first time an extension of the profile to support architectural styles [18] and dynamic reconfiguration.

The formal semantics of the extension is defined in terms of *Architectural Design Rewriting* (ADR) [6], a graph-based approach for style-based design and reconfiguration of software architectures. This formalism is based on term rewriting and supports

---

<sup>★</sup> This work has been supported by the EU FET-GC2 IP project *SENSORIA*, IST-2005-016004.



**Fig. 1.** The *On road connection* scenario

both style-preserving and style-changing reconfiguration of service-oriented systems. We show how the ADR formalisation and its prototypical implementation in Maude [3], can be used to analyse and specify properties of UML4SOA specifications.

This paper is structured as follows. Section 2 introduces a running example from the automotive domain. Section 3 presents our UML4SOA extension. Section 4 describes the formal ADR semantics of the profile. In Section 5 we describe how our prototypical implementation can be exploited to analyse and verify properties of UML4SOA specifications. Sections 6 and 7 discuss related work and draw some conclusions.

## 2 Running Example

We illustrate our approach by exploiting the *On Road Connectivity* scenario from a SENSORIA case study where a road assistance group of services support car drivers activities. In the scenario, cars access wireless services via stations that are situated along a road. We use a UML component to represent a configuration of such a system. Figure 1 shows the white-box view of a system as a component (top-left box) that contains other components (nested boxes) as parts of its internal assembly. A car is connected to the service access point of a station, which can be shared with other cars that are attached to the same station. A station and its accessing cars form a *cell*, which is dynamically reconfigurable, in the sense that cars can move away from the range of the station of their current cell and enter the range of another cell. A handover protocol permits cars to migrate to adjacent cells as in standard cellular networks.

Stations, in addition to the service access point, use two other communication ports that we call chaining ports. Such ports are used to link cells in larger *cell chains*. Stations can shut down, in which case their orphan cars are connected to other stations. This is tackled by appropriate system reconfigurations. We shall consider a shut down situation in which orphan cars switch from their normal mode of operation to a cell mode, in which case they become standalone ad-hoc stations (see the CarCells in Fig. 1).

### 3 UML4SOA Extension

UML4SOA [14] is a UML profile that aims to ease the work of software engineers when designing and implementing service-oriented software. UML4SOA is defined as a conservative extension of the UML2 metamodel built on top of the Meta Object Facility (MOF) metamodel, with new elements created as UML stereotypes, tagged values and constraints defined in the Object Constraint Language (OCL). Such a MOF metamodel is the basis for the specification of a model-driven approach for the automated generation of service-oriented software through model transformations. In fact, we defined transformation mechanisms from UML4SOA models to various languages (e.g. BPEL/WSDL [15]). UML4SOA uses extended internal structure and deployment diagrams. The extension for structure diagrams comprises service, service interface and service description [14]. A component may publish several services implemented as ports, which are described by service descriptions. Each service may contain a required and a provided interface that may contain operations. These operations, in turn, may contain an arbitrary number of parameters. The orchestration of these services define a new service. The extension for deployment diagrams is restricted to different types of communication paths between the nodes of a distributed system, i.e. permanent, temporary and on-the-fly [20].

When modelling service oriented applications with our UML4SOA profile we observed the need for convenient mechanisms to model the inherent dynamic topologies of such applications: components join and leave the system and connections are rearranged. Such dynamic reconfigurations provide a number of beneficial features, but require a suitable mechanism to constrain the possible evolutions of system configurations and to avoid ill-formed configurations. In order to express such constraints on topologies, it is common practice to use *architectural styles* [18], i.e. sets of rules specifying the legal constituents of a system configuration and the permitted interconnections between them. Unfortunately, UML offers a limited and unsatisfactory support for architectural styles. We propose a novel extension of UML4SOA to remedy this. In addition we provide a methodology for modelling dynamic changes of configurations under architectural styles.

#### 3.1 The UML4SOA Reconfigurations Profile

We present our UML4SOA extension to draw easily understandable diagrams for architectural styles and reconfigurations. UML4SOA models services with ports. For an enhanced readability we will omit in this paper the «service» stereotype on ports, as all ports in the following represent services. Service providers are represented by components, while (one-to-one) connectors are used to model service references.

*Modelling System Configurations.* We model system configurations with internal structure diagrams. Using such diagrams allows us to model services as ports of service providers. Software engineers can use them to model the internal structure of structured classifiers (such as components), and depict the wirings of fields as well as their names, types, and multiplicities. Service engineers must use the typing and multiplicity features to ensure that the set of instances is restricted to structurally indistinguishable

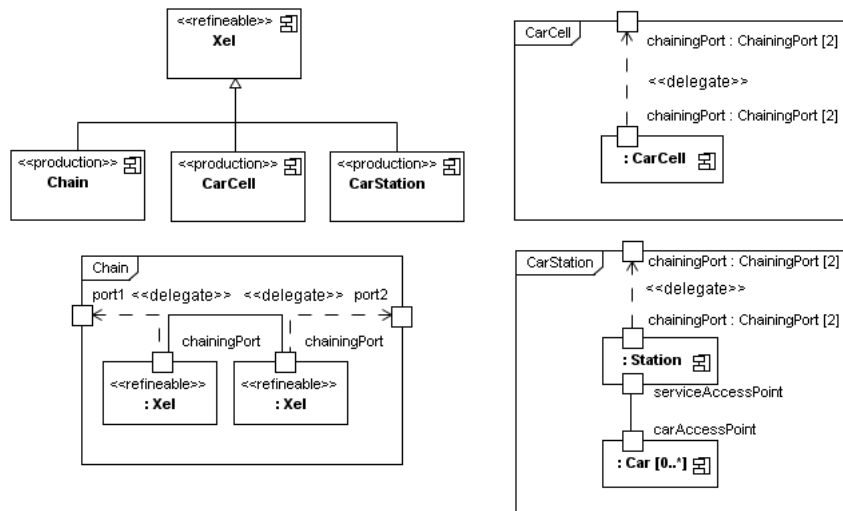


Fig. 2. Architectural style productions for the *On Road Connectivity* scenario

configurations by setting all field multiplicities to a fixed value, and specifying the types of all fields. In this way, instances may only vary with respect to the internal state of the drawn components, but neither in their number nor in their interconnections.

To simplify the modelling of system configurations, we hence introduce the stereotype `<<fragment>>`, which may be applied to components and requires that all component fields (the elements pictured in the internal structure diagram) are typed with component or connector types, and interpret unspecified multiplicities as one. At the same time, we constrain `<<fragment>>` components by forbidding the use range or `*` multiplicities.

Consider Fig. 1 which shows a sample system configuration of our scenario. The internal structure diagram shows constraints on possible interconnection of subcomponents. Obviously (interpreting default multiplicity as one), all instances of this diagram possess the same structure.

A further interesting aspect of internal structure diagrams is that the ports of the container classifier may be drawn on the border of the diagram; if the container classifier owns ports, they may be connected to internal elements by `<<delegates>>` edges, stating that the port of the container classifier represents a proxy of the attached internal port. This allows to define named docking points for the instance, to which other instances may be glued to. For `<<fragment>>` components, we therefore require that all ports must be a delegate of a port of its contained components. As the name implies, a `<<fragment>>` represents a fragment of a system configuration that can be plugged together with other fragments. A complete system configuration is then modeled as a fragment without ports.

*Modelling Architectural Styles.* UML internal structure diagrams provide a set of features to specify architectural styles. Indeed, such diagrams describe the static structure

on the level of types, and allow to constrain multiplicities as well as interconnections. However, the service engineer is forced to model all alternatives allowed by a specific architectural style within one single diagram because of two reasons: first, there is no possibility to define abstract UML components, and second, subtype polymorphism, while present in UML, is not prominently used. We believe that such mechanisms are not enough for a convenient specification of styles.

Our approach, instead, is based on a straightforward extension of the modelling of fragments with two modifications. First, the constraints on multiplicities and on typing of fields are removed. Secondly, to define architectural styles in an inductive manner with composable patterns, by using «refineable» components instead of concrete ones. Components used to define architectural style patterns are tagged with the stereotype «production». The non-terminal components marked with «refineable» may be replaced by any specializing «production» pattern. In our scenario, for example, the «production» Chain (cf. Fig. 2) contains two occurrences of «refineable» Xel components, which may be replaced by Chain, CarCell, and CarStation and productions, as they all specialize the «refineable» Xel.

The «production» patterns define the legal wirings between components, and at the same time represent the basic building blocks of an architectural style. An architectural style is represented by a set of «production» patterns in the sense that every legal configuration must be produceable by applying the production patterns of the architectural style, replacing «refineable» by specializing «production» patterns.

Similarly, it is possible to analyse a given system configuration with the help of production rules, and to determine whether it adheres to the architectural style defined by the used production rules.

*Modeling Reconfigurations under Architectural Styles.* Reconfiguration rules are defined as «transformation» packages having two «pattern» stereotyped components with internal structure diagrams (a left hand side and right hand side pattern), linked with a «transforms» edge. The name of fields in «pattern» components is interpreted as variable names and the enclosing «transformation» package as scope, hence allowing to share variables among left hand side and right hand side patterns.

Often enough, reconfiguration rules depends on complex or non-local conditions. Consider for example the shut-down of a connection station: The connected cars should form an ad-hoc network chain which will be connected to the neighbouring operating stations. Having only simple rules at hand, one would have to write one rule for each possible number of cars to be reconfigured. Using recursive rules, that is to say, using application conditions as in conditional rewrite frameworks, allows us to model the reconfiguration of arbitrary many cars to a linear ad-hoc network, as in Fig. 3.

As can be seen from the diagram, one stereotype was introduced, «preconditions», which is attached to a dependency edge and points to a package containing the reconfiguration preconditions. The scope of variables in precondition patterns is again the enclosing «transformation» package, hence allowing reconfigured architectures to carry over transformation results from the preconditions to the actually performed reconfiguration step. In this way, a complex reconfiguration involving arbitrary many components may be modelled using simple and local reconfigurations.

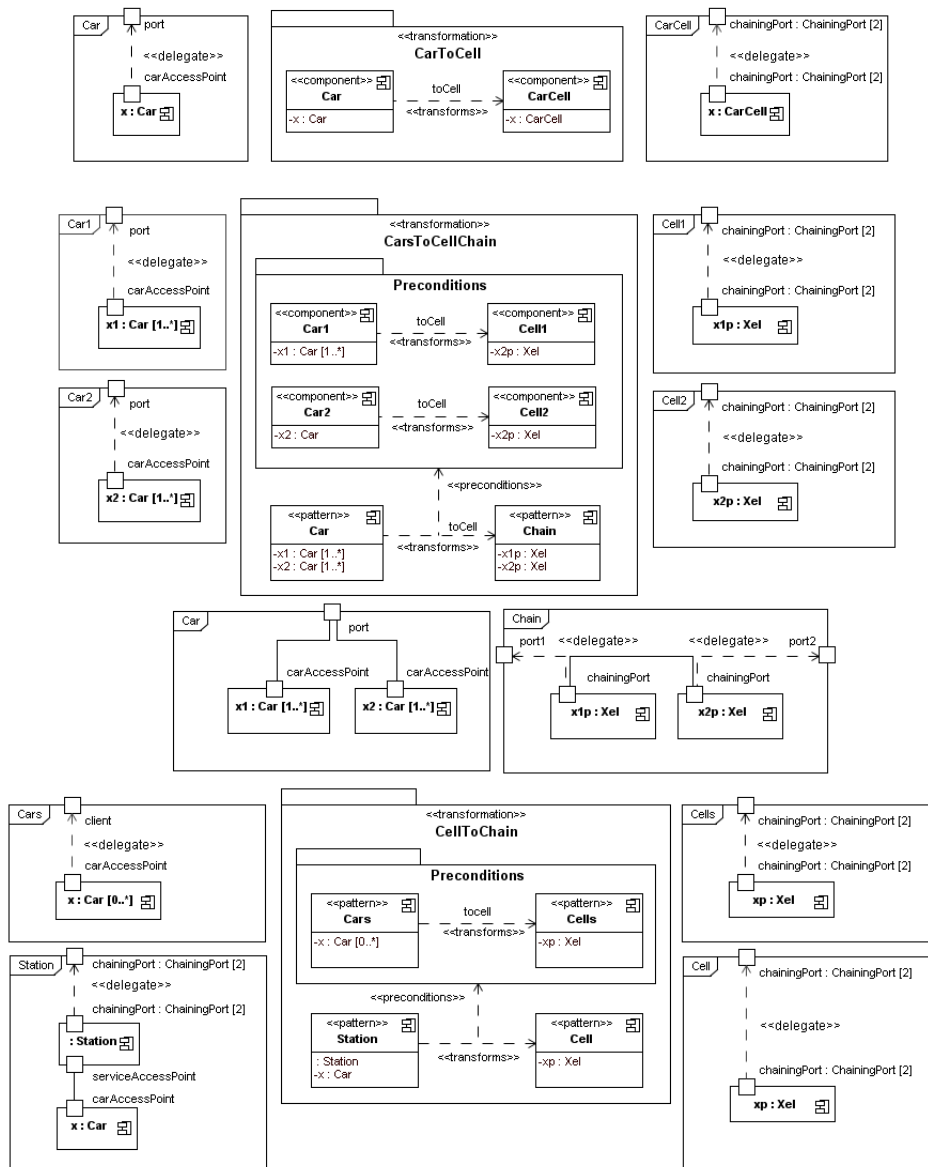


Fig. 3. Ad-hoc network reconfiguration rules for the *On Road Connectivity* scenario

One major challenge when modelling dynamic reconfigurations is to guarantee that the constraints of the architectural style are not violated. The benefit of our approach is that style preservation is ensured just by having the same types in the left- and right-hand sides of reconfiguration rules.

Overall, the UML4SOA profile for architectural design allows to model sample instances of architectural styles, production rules that define an architectural style and that can be composed to create system configurations, and a declarative rule-based framework that allows to model complex reconfigurations of system configurations.

## 4 Formal support


ADR [6] models systems by *designs*: a kind of typed, interfaced graphs whose inner items represent the system components and their interconnections and whose interfaces express their overall types and connection capabilities. Domains of *valid* systems, (e.g. those compliant to styles) are defined in an inductive way by means of *design productions* (i.e. valid system compositions), which define an algebra of *design terms*, each encoding the structure of the system and providing a proof of validity (e.g. style conformance). Reconfiguration and behaviour are given as term rewrite rules acting over design terms rather than over designs. This enables the flexible definition of valid (e.g. style preserving) reconfigurations. A prototypical implementation is described in [3], where we also extended the approach to the treatment of hierarchical graphs and explained how to write system specifications and how to analyse them. ADR has been already validated over heterogeneous models such as network topologies, architectural styles and modelling languages. For instance, in [4] we presented a formalisation of design and reconfiguration aspects of SRML, SENSORIA's business-level service modelling language.

### 4.1 ADR Semantics for the UML4SOA Reconfiguration Profile

This section describes, in an illustrative manner, the ADR formal semantics of the above presented UML4SOA profile. The main idea behind the formalisation is that «fragment»-stereotyped components, i.e. configurations, are represented by ADR designs, while the architectural constraints imposed by UML concepts such as multiplicity or productions are captured by appropriate ADR types and design productions. UML4SOA reconfiguration rules specified as «transformation» packages are represented by ADR rewrite rules. It is worth to recall that the main novel principles of the profile, i.e. style-consistent design-by-refinement and style-preserving, conditional reconfigurations are indeed the quintessence of ADR.

*Modelling System Configurations in ADR.* A design is a graph-based structure. Recall that a *graph* is a tuple  $G = \langle V, E, \theta \rangle$  where  $V$  is the set of nodes,  $E$  is the set of edges and  $\theta : E \rightarrow V^*$  is the tentacle function. Given a graph  $T$  (called the *type graph*), a *T-typed graph* is a pair  $\langle G, t_G : G \rightarrow T \rangle$ , where  $G$  is the *underlying graph* and  $t_G : G \rightarrow T$  is a graph morphism. From now on we assume that graphs are *T*-typed.

Technically, a *design* is a triple  $d = \langle L_d, R_d, i_d \rangle$ , where  $L_d$  is the interface graph consisting of a single so-called *non-terminal edge* (the *interface*) whose tentacles are attached to distinct nodes;  $R_d$  is the body graph; and  $i_d : V_{L_d} \rightarrow V_{R_d}$  is the total function that maps interface nodes to body nodes.



**Fig. 4.** The *On Road Connectivity* scenario of Fig. 1 as an ADR design

The visual representation of a design (see Fig. 4) depicts the interface as a dotted box with its type written in its top-left corner. The body is depicted inside the dotted box. Edges are represented as boxes (possibly rounded), tentacles as arrows (their order is given by their orientation) and nodes as small circles. The nodes being exposed on the interface are denoted by wavy lines.

Fig. 4 exemplifies how UML4SOA *«fragment»* components of Fig. 1 can be mapped to ADR designs: *«service»* ports are mapped to ADR nodes, while the port type determines the node type (e.g. UML types *ChainingPort*, *CarAccessPort* and *StationAccessPort* are represented by node types  $\bullet$ ,  $\circ$  and  $\odot$ , respectively). Components are mapped to hyper-edges, where the component type determines the hyper-edge type. UML connectors are mapped to binary edges of a predefined type *c*.


The interface of the design is defined by the ports and the generalisation of the *«fragment»* component. The ports of the *«fragment»* define the set of interface nodes  $V_{L_d}$ , and each *«delegates»* edge defines a maplet of the mapping  $i_d$  from interface to body nodes  $V_{R_d}$ . The type of the so-produced graph, as defined by the UML4SOA model, is determined by the generalisation of each *«fragment»* (*Xel* in Fig. 1).

*Modelling Architectural Styles in ADR.* The distinction between refinable components and non-refinable components amounts to the distinction between non-terminal and terminal edges in ADR. The underlying idea is the same: a non-terminal edge is an edge intended to be refined (i.e. replaced by an arbitrarily complex graph). Non-terminal edges can appear in designs, representing unspecified parts of a configuration (a refinable component) or in design productions (see later). Terminal edges instead represent parts of a graph that cannot be further refined (non-refinable components).

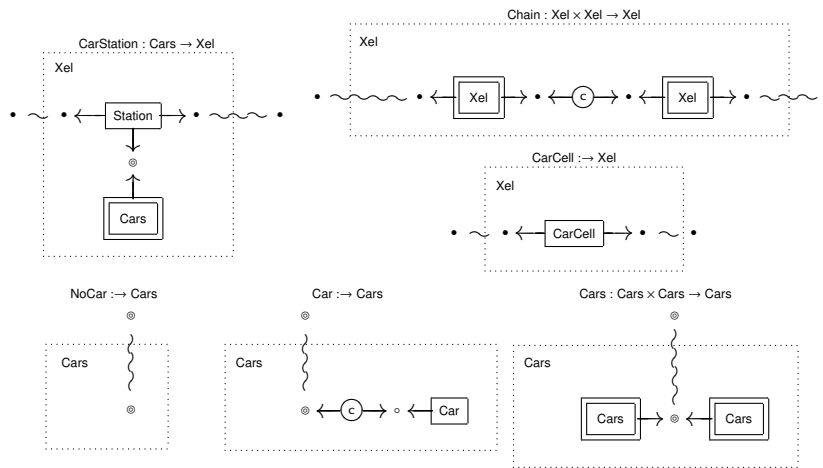
The style definition mechanisms of UML4SOA, i.e. internal structure diagrams and productions, are modelled by ADR type graphs and by design productions. Note however that some of the architectural constraints involved in class diagrams such as multiplicities cannot be directly mimicked by type graphs. Instead, they are dealt with at the level of design productions.

Consider the type graph of left on Fig. 5. It is easy to see that each edge corresponds to a component (*Car*, *Station*, *CarCell*) or connector type (the overloaded symbol *c*).






**Fig. 5.** Type graph for *On Road Connectivity* scenario (left) and a correctly typed graph(right)



**Fig. 6.** Design productions for *On Road Connectivity* scenario

Non-terminal edges (Cars, Xel) are distinguished by their double border. In general, the type graph is obtained from the whole UML4SOA specification: adding terminal edges for each non-refinable component type, non-terminal edges for each refinable component type, nodes for port types, tentacles for component ports and edges for the connectors.

Type graphs do not impose any multiplicity constraint, i.e. they would amount to a UML [0..\*] multiplicity constraint. A suitable way to impose a multiplicity constraint in ADR is by means of design productions. For instance, the treatment of sets of cars in the UML4SOA specification via multiplicities is dealt in ADR with the design productions NoCar, Car and Cars (see Fig. 6), which respectively allow to refine a generic set of cars as an empty set, a singleton or the union of two other sets. In this way, UML4SOA productions are directly mapped into ADR design productions. For instance, in absence of production NoCar the multiplicity constraint would be [1..\*]. We remark that productions allow to refine the architectural constraints imposed by a type graph alone. For instance, the graph on the right of Fig. 5 is well-typed but is not generated by our productions.



**Fig. 7.** Reconfiguration CarToCell

Technically, a design production is very much like a design but with an order on the non-terminal edges of the body graph (intuitively, the order of the arguments they represent). The *type* of a production  $p$  is  $A_1 \times A_2 \times \dots \times A_{n_p} \rightarrow A_p$ , where  $A_k$  is the non-terminal symbol labelling the  $k$ -th non-terminal edge  $e_k$  of the body of the production. The functional type  $A_1 \times A_2 \times \dots \times A_{n_p} \rightarrow A_p$  associated to a production  $p$  is not an accident. In fact,  $p$  can be considered a function that when applied to a tuple  $\langle d_1, d_2, \dots, d_{n_p} \rangle$  of designs of types  $A_1, A_2, \dots, A_{n_p}$ , respectively, returns a design  $d = p(d_1, d_2, \dots, d_{n_p})$  of type  $A_p$ . The definition is obvious:  $d = (L_p, R_d, i_p)$ , where  $R_d$  is obtained from  $R_p$  by replacing edge  $e_k$  in it with graph  $R_{d_k}$  respecting the tentacle function  $i_{d_k}$ ,  $k = 1, \dots, n_p$ .

This view corresponds to a bottom-up design development: a design is constructed by putting together some component designs. However, the dual view is also possible: a production can be seen as a refinement of an abstract component of type  $A$  as an assembly of concrete and abstract components, the latter being of type  $A_1, A_2, \dots, A_{n_p}$ .

*Modelling Reconfigurations under Architectural Styles.* UML4SOA transformations are represented by ADR rewrite rules. We just recall here that one of the advantages of ADR reconfigurations over other graph-based approaches is style-preservation, which is guaranteed by rewrites that do not change the overall type (they can actually change the type of certain sub-parts in the rule derivation of the overall reconfiguration).

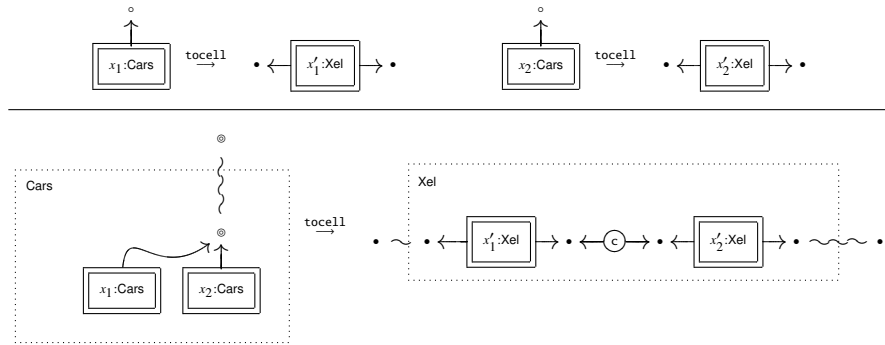
Translating UML4SOA reconfiguration rules to ADR in the general case is done by UML4SOA preconditions, «transforms» left-hand and right hand sides, and transformation labels are translated to their respective counterparts in ADR. In this process, «pattern» components are translated to ADR designs by first producing ADR design graphs (replacing components with  $[0..*]$  multiplicities by the corresponding non-terminal hyper-edge, as done in the example with Cars) and then parsing the result using the ADR productions generated from the UML4SOA productions.

The ad-hoc network reconfiguration is tackled by using inductive reconfiguration rules in SOS style. The base reconfiguration involves a single car (see Fig. 7):

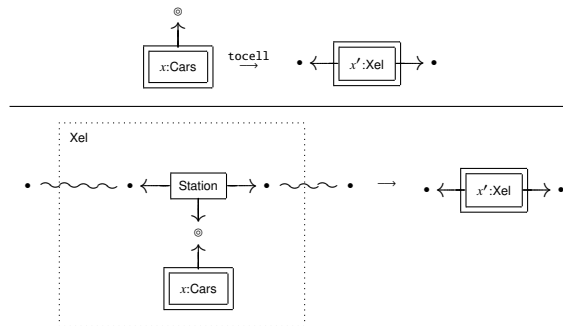
$$\text{CarToCell} : \text{Car} \xrightarrow{\text{toCell}} \text{CarCell}$$

The inductive case we consider is illustrated in Fig. 8, where the union of two collections of cars is reconfigured as the concatenation of the respective reconfigured cells, provided that these are possible:

$$\text{CarsToCellChain} : \frac{x_1 \xrightarrow{\text{toCell}} x'_1 \quad x_2 \xrightarrow{\text{toCell}} x'_2}{\text{Cars}(x_1, x_2) \xrightarrow{\text{toCell}} \text{Chain}(x'_1, x'_2)}$$



**Fig. 8.** Reconfiguration CarsToCellChain



**Fig. 9.** Reconfiguration CellToChain

The cell with the station shutting down is reconfigured by the rule (see Fig. 9):

$$\text{CellToChain} : \frac{x \xrightarrow{\text{toCell}} x'}{\text{CarStation}(x) \longrightarrow x'}$$

Obviously, types are not preserved by CarToCell and CarsToCellChain and thus the right- and left-hand sides of the rewriting rule cannot be applied in the same contexts. Type changing allows for the modelling of reconfigurations that lead from one architectural style to another. However, this is not what we want in this example and thus labelled rules are given in SOS style. The last rule CellToChain, instead, is given as a conditional term rewrite rule, where the premise is for a collection of cars to become a chain cell, while the conclusion actually transforms a chain of cells into a chain of cells. The type is preserved and the silent label makes it applicable in any larger context (unlike style-changing rewrites labelled toCell).

## 5 Analysis and Verification

This section emphasizes the benefit of having a formal semantics for our UML profile by describing how the use of our implementation of ADR [3] can be used to analyse and verify properties of UML4SOA specifications. We remark that the implementation of the formalisation, i.e. the translation of UML4SOA specifications to ADR specifications has not been implemented yet. On the other hand, a prototypical implementation of ADR is available for download [3]. Nevertheless, we offer sufficient evidence of the potential of our approach as a helpful support for software architects.

*Analysing Styles.* After a first development of a UML4SOA specification, a software architect might wonder whether the defined architectural styles enjoy some properties he desires. For instance, in our example scenario one could be interested in stating that no Xel production builds a configuration in which the left and right chaining ports are disconnected. It is easy to see that this property trivially holds in our example. However, as scenarios become complicated such properties become more subtle. Note that due to the inductive definition of styles it holds that if all Xel productions satisfy the property, then the property holds for any possible Xel configuration. This is indeed the case of the example property: all Xels are un-broken chains.

Our implementation includes a graph logic (Courcelle's MSO) that allow us to reason about the structure of a graph. Such mechanism can be used to analyse structure of UML4SOA productions by analysing the underlying graphs. The above example for instance is a well known property of graph connectivity which can be expressed in MSO by  $\forall X.((\forall x, y(y \in X \wedge z \in R(y, z) \rightarrow z \in X \wedge \forall y.R(a, y) \rightarrow y \in X)) \rightarrow b \in X)$ , where  $X$  is a set of nodes,  $x, y$  and  $z$  are nodes,  $R$  abbreviates the existence of an edge between two nodes, and  $a, b$  are shorthands for the left and right hand-side chaining ports. In words, we look for a all sets of nodes  $X$  closed under the transitive closure of the relation  $R$  of direct adjacency and containing all nodes adjacent to  $b$ . If all such sets contain  $b$  too, then  $a$  and  $b$  reachable from each other. The above formula holds for all body graphs of Xel productions. We can of course, write abbreviations for such formulae to construct a sort of library of structural properties.

*Checking Style Conformance.* Once the software architect is confident with the style he has designed he might be interested in re-using some of his old specifications. After manually applying some cosmetics on the types of diagrams and other entities, he might want to know whether the resulting instance is consistent with the style. Roughly, he needs a correct parsing in terms of the productions. This is supported by a mechanism that roughly generates design terms and checks if the resulting designs are isomorphic to the configuration under analysis. For instance, one can show that the configuration in Fig. 4 is style conformant by finding the parsing that we mentioned in Section 3.1. A counterexample can be found in Fig. 5 (right).

*Finding Configurations Automatically.* Model finding is the problem of analysing the state space of all possible instances of his architectural style. Such analysis serves as a computer-aided design process or as a debugging method to find out inconsistencies in models, styles or properties. Our model finding system is based on two mechanisms:

one to generate a state space of models and one to explore it. In our approach we can define a rewrite theory that simulates a design-by-refinement process, roughly consisting of the context-free graph grammar obtained by a left-to-right reading of design productions. In order to explore such state spaces we can use various mechanisms of Maude. Typically, the space of configurations is infinite and bounds are required. For instance, we can use search strategies to find configurations with at 4 cars, 3 stations and 2 car cells and we obtain, among others, the design of Fig. 4.

*Analysing Static Aspects of Configurations.* Now that the software architect has adapted some of his old designs and possibly built new ones, he might want to reason about them. Returning to our example, we might wonder if a configuration has at least  $n$  cars or is free of cars in cell mode. Recall that our configurations have two levels: the more abstract level of design terms and the more detailed level of the diagrams. We can expect dual mechanisms for stating structural aspects. Indeed, we saw above that the properties of diagrams are supported by graph logics. Similarly, properties can be stated at the level of design terms. Our ADR implementation does this by means of spatial logics, the natural and structured way to reason about term-like specifications. Basically, for each design production  $f$  used to compose designs the logic incorporates a spatial operator  $f$ -so to decompose a design. For instance, formula  $\text{Chain-so}(\phi_1, \phi_2)$  is satisfied by all those designs of the form  $\text{Chain}(x_1, x_2)$ , where design  $x_1$  satisfies formula  $\phi_1$  and design  $x_2$  satisfies formula  $\phi_2$ .

Consider the property that states a collection of cars has at least  $n$  cars. We can inductively define it as follows: For  $n$  equal to zero the formula always holds. For  $n + 1$  the formula holds whenever the term is decomposable as the composition via operation  $\text{Cars}$  of one car ( $\text{Car-so}$ ) and a term with at least  $n$  cars. Using such formulas we can for instance check that the design of Fig. 4 satisfies the property stating that each station has at least one car and violates the property stating that each station has at least two cars.

*Analysing Dynamic Instances of Configurations.* At this point the software architect might be confident with the structural properties enjoyed by his configurations. The modelled application, however, has a dynamic architecture with a various reconfiguration rules as those we use in our running example. Can he express that some property is invariantly preserved or that some bad property will never happen? The standard way to reason about such properties is by means of temporal logics.

In our case temporal logics are supported Maude's built-in LTL model checker. Properties regarding dynamic aspects of reconfigurations are expressed using the Linear-time Temporal Logic (LTL). Roughly, one is able to reason about infinite sequences of reconfigurations, by expressing properties on the ordering of state (i.e. configuration) observations. Such observations are predicates expressing structural properties as above mentioned.

As an example we can write the formula asserting that *it is always true that a collection of cars has at least 2 cars* as  $\Box \text{at-least-k-cars}(2)$ , where  $\Box$  denotes the *always* temporal operator. This property trivially holds for the design of figure 4. Indeed, no reconfiguration rule allows cars to leave the system so that their number remains constant.

## 6 Related Work

The Service Component Architecture (SCA) [8] focuses on policies and implementation aspects of services but is not based on UML. The work in [19] is based on UML models and transformations to executable descriptions of services. However, the approach lacks an appropriate UML profile preventing one from building models at the high level of abstraction; thus producing overloaded diagrams. The work of [11] proposes to use modes to address dynamic reconfiguration of service-oriented architectures and extends the UML to visualize such reconfiguration. The UML extension sticks to the mode terminology and does not include a visualization of the transformation rules. The OMG is also working to standardize a UML profile and metamodel for services (UPMS) [17]. The current version does not support styles or reconfigurations.

Structural aspects for services, modelled in UML, have also been addressed in several other works (e.g. [13]); however, as far as we know, none of them is based on a formal background like the one presented here. The only exception is the UML extension for service-oriented architectures that can be found in [2]. The approach includes refinement issues based on architectural styles and is formalized by graph transformation systems. It includes stereotypes for the structural specification of services. However, it does not introduce specific model elements for the orchestration of services, the notion of style there is less expressive (it basically amounts to our type graphs) and reconfigurations there are limited to unconditional ones.

A completely different approach to modelling architectural styles in UML would be to use constraints expressed in the OMG Object Constraint Language (OCL). To the best of our knowledge, however, there is no reconfiguration approach using solely OCL, which would have two drawbacks: OCL is a textual notation and it would introduce another language to the service engineer.

ADR has been mainly inspired by graph-based approaches to architectural styles [12,16] (see [6] for a comparison). The use of graphs and graph transformations to model architectural styles has been proposed by several authors (see [18], for instance) who based their approaches on the concept of *shapes* in programming languages. ADR shares also concepts with approaches based on process calculi with reconfigurable components (e.g. [1]). The main advantages of ADR are that the hierarchical and inductively based approach allows us to compactly represent complex reconfiguration rules, and that, style preservation is guaranteed by construction. ADR is also related to approaches that deal with reconfigurations in software architectures defined by an ADL (see [5]). The main advantages of ADR are the unified treatment of design, behaviours and reconfiguration, and the use of hierarchical, inductive reconfigurations. A comparison with a logic based architectural design methodology was given in [7].

## 7 Conclusion

We have presented a novel extension of UML4SOA, our approach for the modelling of service oriented architectures. The profile offers suitable ingredients to deal with architectural styles and reconfigurations. We have equipped the proposed profile with a formal semantics, offering support for analysis and verification from the very early

stages of modeling. Thus, our approach is a comprehensive and pragmatic but theoretically well founded approach to software engineering for service-oriented systems. Our current efforts are aimed completing our tool support by automatising the translation of UML4SOA specifications and upgrading the prototypical implementation of ADR to a real tool. In future work we would like to integrate our approach in the SENSORIA suite of tools and techniques, which already includes some development [10] and re-engineering (legacy systems as services) [9] instruments.

## References

1. N. Aguirre and T. S. E. Maibaum. Hierarchical temporal specifications of dynamically reconfigurable component based systems. *ENTCS*, 108:69–81, 2004.
2. L. Baresi, R. Heckel, S. Thöne, and D. Varró. Style-based modeling and refinement of service-oriented architectures. *SOSYM*, 5(2):187–207, 2006.
3. R. Bruni, A. Lluch Lafuente, and U. Montanari. Hierarchical Design Rewriting with Maude. In *WRLA'08*, ENTCS. Elsevier, 2008. To appear.
4. R. Bruni, A. Lluch Lafuente, U. Montanari, and Emilio Tuosto. Service Oriented Architectural Design. In *TGC'07*, volume 4912 of *LNCS*, pages 186–203. Springer, 2008.
5. R. Bruni, A. Lluch Lafuente, U. Montanari, and E. Tuosto. Architectural Design Rewriting as an Architecture Description Language. *R2D2* Microsoft Research Meeting, 2008.
6. R. Bruni, A. Lluch Lafuente, U. Montanari, and E. Tuosto. Style Based Architectural Reconfigurations. In *EATCS Bulletin*, volume 94, pages 161–180. February 2008.
7. A. Bucchiarone, R. Bruni, S. Gnesi, and A. Lluch Lafuente. Graph-Based Design and Analysis of Dynamic Software Architectures. In *Concurrency, Graphs and Models*, volume 5065 of *LNCS*. Springer, 2008.
8. S. Consortium. Service Component Architecture Policy Framework, Version 1.0, 2007.
9. R. Correia, C. Matos, R. Heckel, and M. El-Ramly. Architecture migration driven by code categorization. In *ECSA*, volume 4758 of *LNCS*, pages 115–122. Springer, 2007.
10. H. Foster and P. Mayer. Leveraging integrated tools for model-based analysis of service compositions. In *ICIW'08*. IEEE Computer Society Press, 2008.
11. H. Foster, S. Uchitel, J. Kramer, and J. Magee. Leveraging Modes and UML2 for Service Brokering Specifications. volume 389 of *LNCS*, pages 76–90. CEUR, 2008.
12. D. Hirsch and U. Montanari. Shaped hierarchical architectural design. *ENTCS*, 109, 2004.
13. S. Johnson. UML 2.0 Profile for Software Services, 2005.
14. N. Koch, P. Mayer, R. Heckel, L. Göczi, and C. Montangero. D1.4a: UML for Service-Oriented Systems. Specification, SENSORIA Project 016004, 2007.
15. P. Mayer, A. Schroeder, and N. Koch. A Model-Driven Approach to Service Orchestration. In *SCC'08*, IEEE, pages 1–6. IEEE, 2008.
16. D. L. Métayer. Describing software architecture styles using graph grammars. *IEEE Transactions on Software Engineering*, 24(7):521–533, 1998.
17. Object Management Group (OMG). UML Profile and Metamodel for Services, 2008.
18. M. Shaw and D. Garlan. *Software Architecture: Perspectives on an Emerging Discipline*. Prentice-Hall, New Jersey, USA, 1996.
19. D. Skogan, R. Grønmo, and I. Solheim. Web service composition in UML. In *EDOC'04*, pages 47–57. IEEE Computer Society, 2004.
20. M. Wirsing, A. Clark, S. Gilmore, M. Hölzl, A. Knapp, N. Koch, and A. Schroeder. Semantic-Based Development of Service-Oriented Systems. In *FORTE'06*, volume 4229 of *LNCS*, pages 24–45. Springer, 2006.