Modelling Dynamic Software Architectures using Typed Graph Grammars [1]

Roberto Bruni [a], Antonio Bucchiarone [b,2], Stefania Gnesi [c] and Hernán Melgratti [b]

[a] *Dipartimento di Informatica, Università di Pisa, Italy. Email: bruni@di.unipi.it*

[b] *IMT Alti Studi Lucca, Italy. Email: {a.bucchiarone,h.melgratti}@imtlucca.it*

[c] *ISTI-CNR, Pisa, Italy. Email: stefania.gnesi@isti.cnr.it*

Abstract

Several recent research efforts have focused on the dynamic aspects of software architectures providing suitable models and techniques for handling the run-time modification of the structure of a system. A large number of heterogeneous proposals for addressing dynamic architectures at many different levels of abstraction have been provided, such as programmable, ad-hoc, self-healing and self-repairing among others. It is then important to have a clear picture of the relations among these proposals by formulating them into a uniform framework and contrasting the different verification aspects that can be reasonably addressed by each proposal. Our work is a contribution in this line. In particular, we map several notions of dynamicity into the same formal framework in order to distill the similarities and differences among them. As a result we explain different styles of architectural dynamisms in term of graph grammars and get some better insights on the kinds of formal properties that can be naturally associated to such different specification styles. We take a simple automotive scenario as a running example to illustrate main ideas.

Key words: Dynamic Software Architectures, Typed Graph Grammars and Modelling.

1 Introduction

In the last decades, computer systems have changed from isolated static devices to highly interconnected machines that execute their tasks in a cooperative and

[1] Research supported by the EU within the FET-GC2 IST-2005-16004 Integrated Project SENSORIA (*Software Engineering for Service-Oriented Overlay Computers*) and by the Italian FIRB Project TOCAI.IT.

[2] Also supported by the Marie Curie Host Fellowships for Transfer of Knowledge (FP6-2002-Mobility 3 Proposal n. FP6-14525) and by Nokia Siemens Networks, Lisboa, Portugal.

coordinated manner. These modern, complex systems are known as *global computing systems* (GCS) or *network-aware computers*, and have to deal with frequent changes of the network environment. In a GCS, components are autonomous and dynamic, the network's coverage is variable, and there is not a centralized authority.

Software architectural models are intended to describe the structure of a system in terms of computational components, their interactions, and its composition patterns [23], so to reason about systems at a more abstract level, disregarding implementation details. Since GCS may change at design time, pre-execution time, or run-time [20], software architecture models for GCS should be able to describe the changes of the system structure and to enact the modifications during the system execution [19]. Such models are generally referred to as *Dynamic Software Architectures (*DSA*s)*, to emphasize that the system architecture evolves during runtime.

A variety of definitions of dynamicity for software architecture have been proposed in the literature. Below we list some of the most prominent definitions to show the variability of connotations that the word *dynamic* acquires.

- **Programmed Dynamism** [8]. All admissible changes are defined prior to run-time and are triggered by the system itself.
- **Self-repairing** [22]. Changes are initiated and assessed internally, i.e., the run-time behavior of the system is monitored to determine whether a change is needed. In such case, a reconfiguration is automatically performed.
- **Self-adaptive** [21]. Systems can adapt to their environments by enacting run-time changes.
- **Ad-hoc dynamism** [8]. Changes are initiated by the user as part of a software maintenance task, they are defined at run-time and are not known at design-time.
- **Constructible dynamism** [2]. It is a kind of ad-hoc mechanism but all architectural changes must be described in a given modification language, whose primitives constrain the admissible changes.

The different proposals for DSA are bound to particular languages and models. In this paper we are aimed at understanding the main notions relying behind such proposals by abstracting away from particular languages and notations. We want to give a uniform formal presentation that is abstract enough to cover most of those features. In this sense, our work is in the line of other previous research efforts [24,7]. In particular we select graph grammars as a formal framework for mapping the different notions of dynamicity because (i) they provide both a formal basis and a graphical representation that is in line with the usual way architectures are represented, (ii) they allows for a natural way of describing styles and configurations, (iii) they have been largely used for specifying architectures. The use of graph grammars is instrumental in comparing different mechanisms and better understanding the kinds of properties that can be naturally associated to such specifications. We argue that the characterisation of dynamicity we present is to some extent orthogonal to the particular kind of graph grammars we use, and therefore, extensible to different variants of graph rewriting systems.

Related Work. Several previous works have proposed alternative ways for describing software architectures by using graph grammar. Our representation of DSA as graph grammars is borrowed from the Le Métayer approach [16]. Actually the notion of programmed DSA corresponds to that proposal. A different way of representing software architectures with graphs can be found in [12], where hyperedges are components and nodes are ports of communication, and the reconfiguration is given as context-free productions together with a constraint solving mechanism. Also Baresi et al. in [3,4] use graph transformation systems to model programmed architectural styles at different levels of abstraction. Self-repairing mechanisms have been proposed in [1,9,10,22]. Ad-hoc reconfiguration has been studied in [8] as a programming language that allows for the runtime modification of software architectures. Similarly, proposals for constructible languages can be found in [20]. These approaches are interested in providing executable frameworks for supporting DSA: The main difference w.r.t. our work is that they are aimed at providing real specification/ programming/ languages while we are aimed at giving an abstract characterization of such kind of mechanisms.

As far as the different flavours of dynamicity are concerned, the work of Wermelinger in [24] explores the ability of the Chemical Abstract Machine (CHAM) [5] to express the dynamics of software architectures. His formalization proposes particular CHAM (and commands) to tackle self-organized, ad-hoc, and programmed reconfiguration. Differently, we are interested in understanding how each particular style of dynamism is reflected into a graph grammar.

Organization. In Section 2 we describe the formal framework used in the rest of the paper, and the way in which software architectures are represented by using hypergraphs. Then we show how different forms of dynamism in software architecture can be expressed in terms of graph grammars (Section 3) and apply them to a simple case study (Section 4). Other orthogonal aspects of dynamism are discussed in Section 5. Some final remarks and future lines of research are in Section 6.

2 Formalization of Dynamicity

We model components and connectors as hyperedges and the ports to which they are attached as nodes. Figure 1 depicts an hypergraph containing two nodes port_1 and port_2, the hyperedge $\mathsf{component}$ (a component that exposes two different ports), and the hyperedge $\mathsf{connector}$ (a connector that has two tentacles to the port port_1 and one to the port port_2). Note that component edges are drawn as square boxes while connector edges as rounded boxes. Moreover, we show the ordering of tentacles by labeling the corresponding arrows with natural numbers (in some cases we shall use suitable names instead of numbers as labels, so to ease the reading).

Definition 2.1 [Hypergraph] A *(hyper)graph* is a triple $H = (N_H, E_H, \phi_H)$, where N_H is the set of nodes, E_H is the set of (hyper)edges, and $\phi_H : E_H \to N_H^+$ describes the connections of the graph, where N_H^+ stands for the set of non-empty strings of elements of N_H. We call $|\phi_H(e)|$ the *rank* of e, with $|\phi_H(e)| > 0$ for any $e \in E_H$.

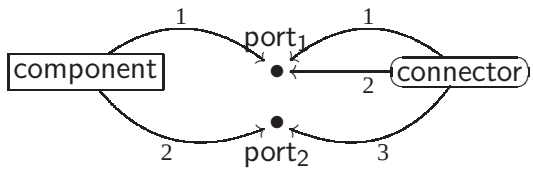

Figure 1. A hypergraph describing a style.

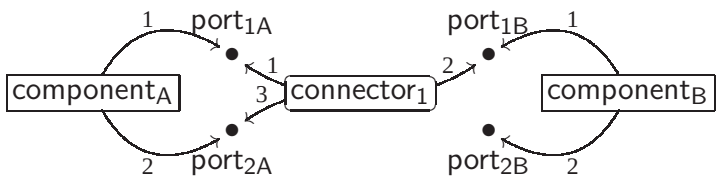

Figure 2. A hypergraph describing a configuration for the style in Figure 1

The connection function ϕ_H associates each hyperedge e to the ordered, non empty sequence of nodes e is attached to. An architectural style is just a hypergraph T that describes only the types of ports, connectors, components and the allowed connections. A configuration compliant to such style is then described by the notion of a T-typed hypergraph.

Definition 2.2 [Typed Hypergraph] Given a hypergraph T (called the *style*), a T-*typed hypergraph* or *configuration* is a pair $\langle |G|, \tau_G \rangle$, where $|G|$ is the *underlying* graph and $\tau_G : |G| \to T$ is a total hypergraph morphism.

The graph $|G|$ defines the configuration of the system, while τ_G defines the (static) *typing* of the resources. We recall that a total hypergraph morphism $f : G \to G'$ is a couple $f = \langle f_N : N \to N', f_E : E \to E' \rangle$ such that: $f_N(\phi_G(e)) = \phi_{G'}(f_E(e))$ (we overload f_N to denote also the homomorphic extension of f_N over strings).

Consider the style T in Figure 1: there is one unique type component of components exposing two ports of different types, and one connector attached to two ports of type port_1 and one port of type port_2. Then, a possible T-typed hypergraph (or a configuration of the style T) is in Figure 2: it has two different components with their corresponding ports, and one connector. The typing morphism is implicitly defined by the name of the elements in the configuration, which consist of the type name plus a subindex identifying the particular instance (e.g., port $\mathsf{port}_{1\mathsf{A}}$ has type port_1). We remark that the typing morphism requires components to have exactly one port of type port_1 and one of type port_2. Similarly, the only connections valid for a connector are those that attach its first two tentacles to ports of type port_1 and the third one to a port of type port_2. All such constraints are enforced by the existence of a typing morphism.

The reconfiguration of a software architecture is described by a set of rewriting productions. Roughly, a production p is a partial, injective morphism of T-typed graphs, i.e., it has the following shape: $p : L \rightarrowtail R$, where L and R are T-typed hypergraphs, called the *left-hand* and the *right-hand side* of the production, respectively. Given a T-typed graph G and a production p, a rewriting of G using p can be informally described as follow: (1) find a (type preserving) match of the

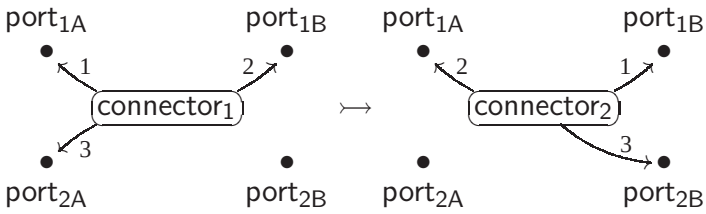

Figure 3. A rewriting production.

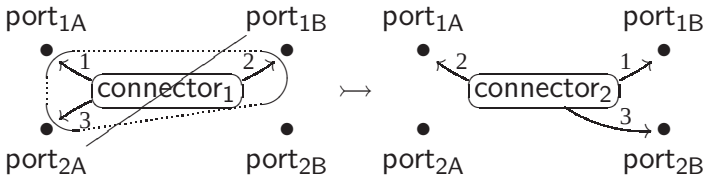

Figure 4. A rewriting production with negative application condition.

left-hand-side L in G, i.e., identify a subgraph of G that corresponds with L; (2) remove from the graph G all the items corresponding to the left-hand side that are not in the right-hand-side; (3) add all the items of the right-hand side that are not in the left-hand-side; (4) the elements that are both in L and R are preserved by the rewrite.

An example of a production is shown in Figure 3 (the morphism among the left and right-hand side of the rule is represented by using the same names for mapped elements, i.e., it is the partial inclusion). The production allows to remove an existing connector $\mathsf{connector}_1$ and to add a new connector $\mathsf{connector}_2$ that is attached to the original ports in a specular way with respect to the original connector.

Finally, an architecture is described by a T-typed graph grammar.

Definition 2.3 [(T-typed) graph grammar] A *(T-typed) graph grammar* $\mathcal{G}$ is a tuple $\langle T, G_{in}, P\rangle$, where G_{in} is the *initial (T-typed) graph* and P is a set of *productions*.

Notation. Let $\mathcal{G} = \langle T, G_{in}, P\rangle$ be a (T-typed) graph grammar, and G and H (T-typed) hypergraphs. We write $G \Rightarrow_p H$ to denote that G is rewritten in one step to H by using the production $p \in P$. We abbreviate the reduction sequence $G_0 \Rightarrow_{p_1} G_1 \Rightarrow_{p_2} \ldots \Rightarrow_{p_n} G_n$ with $G_0 \Rightarrow_{p_1 p_2 \ldots p_n} G_n$. We write $G \Rightarrow^* G'$ to denote that there exists a possible empty sequence $s \in P^*$ of derivation steps such that $G \Rightarrow_s G'$.

For convenience when describing the examples, we will also consider productions with negative application conditions [11], i.e., productions that are equipped with a constraint about the context in which they can be applied. For instance, such conditions can state that the production is applicable only when certain nodes, edges, or subgraphs are not present in the graph. Such conditions are graphically shown in the left-hand-side of a production by grouping forbidden elements into a dotted lined area. Figure 4 shows a production with negative conditions stating that the new connector $\mathsf{connector}_2$ can be added to the configuration if and only if no other connector of type $\mathsf{connector}$ is already attached in a specular way to port_{1A} and port_{1B}. We refer the interested reader to the formal presentation of SPO graph grammars to [17] and to [11] for grammars with negative application conditions.

3 Characterisation of Dynamism

This section characterizes different forms of dynamism in software architecture [2,8,10,20,21,22] in terms of graph grammars. In particular we show that for verification aspects it make sense to focus only on two forms of dynamicity: Programmed and Repairing.

Given a grammar $\mathcal{G} = \langle T, G_{in}, P \rangle$, we will use the following notions:

- The set $\mathcal{R}(\mathcal{G})$ of *reachable configurations*, i.e., all configurations to which the initial configuration G_{in} can evolve. Formally, $\mathcal{R}(\mathcal{G}) = \{G | G_{in} \Rightarrow^* G\}$.
- The set $\mathcal{D}_{\mathsf{P}}(\mathcal{G})$ of *acceptable configurations* of an architecture are defined as the graphs that have type T and satisfies a auitable property P. Formally, $\mathcal{D}_{\mathsf{P}}(\mathcal{G}) = \{G \mid G \textit{ is a } T\textit{-typed graph} \wedge \mathsf{P} \textit{ holds in } G\}$.

3.1 Programmed dynamism

Programmed dynamism assumes that all architectural changes are identified at design time and triggered by the program itself [8]. Many proposals in the literature [16,13,4] that use graph grammars for specifying DSA present this kind of dynamism. A programmed DSA $\mathcal{A}$ is associated with a grammar $\mathcal{G}_{\mathcal{A}} = \langle T, G_{in}, P \rangle$, where T stands for the style of the architecture, G_{in} is the initial configuration, and the set of productions P gives the evolution of the architecture. The grammar fixes the types of all elements in the architecture, and their possible connections, where the productions state the possible ways in which a configuration may change.

Programmed dynamism enables for the formulation of several verification questions. Consider the set of desirable configurations $\mathcal{D}_{\mathsf{P}}(\mathcal{G})$, then it should be possible (at least) to know whether:

- the specification is correct, in the sense that any reachable configuration is desirable. This reduces to prove that $\mathcal{R}(\mathcal{G}) \subseteq \mathcal{D}_{\mathsf{P}}(\mathcal{G})$, or equivalently that $\forall G \in \mathcal{R}(\mathcal{G}) : \mathsf{P} \textit{ holds in } G$.
- the specification is complete, in the sense that any desirable configuration can be reached. This corresponds to prove $\mathcal{D}_{\mathsf{P}}(\mathcal{G}) \subseteq \mathcal{R}(\mathcal{G})$, or equivalently that $\textit{if } \mathsf{P} \textit{ holds in } G \textit{ then } G \in \mathcal{R}(\mathcal{G})$.

Hence, programmed dynamism provides an implicit definition of desirable configurations. That is, the sets of desirable and reachable configurations should coincide, i.e., $\mathcal{D}_p(\mathcal{G}) = \mathcal{R}(\mathcal{G})$.

3.2 Repairing (or healing) dynamism

Self repairing systems are equipped with a mechanism that monitors the system behaviour to determine whether it behaves within prefixed parameters. If a deviation exists, then the system itself is in charge of adapting the configuration [9].

We can think about a repairing architecture as an ordinary graph grammar $\mathcal{G}_{\mathcal{A}} = \langle T, G_{in}, P \rangle$ in which the set of productions is partitioned into three different sets,

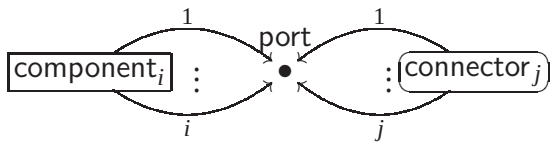

Figure 5. A general graph of types of a software architecture

i.e., $P = P_{pgm} \cup P_{env} \cup P_{rpr}$. Rules in P_{pgm} describe the normal, ideal behaviour of the architecture, i.e., $\mathcal{G}'_{\mathcal{A}} = \langle T, G_{in}, P_{pgm} \rangle$ is a programmed DSA. Rules in P_{env} model the *environment* or, in other words, the ways in which the behaviour of the architecture may deviate from the expected one. Rules in P_{env} may state that the communication among components may be lost or that a non authorised connector become attached to a particular component. Rules P_{rpr} indicate the way in which an undesirable configuration can be repaired in order to become a valid one. That is, the left-hand side of any rule in P_{rpr} identifies a composition pattern in the system that is undesirable. In this way a repairing architecture implicitly defines the desirable configurations of the system as those reachable configurations G that do not exhibit an undesirable composition pattern (i.e., a left-hand-side match for a repairing rule). Formally, the designer would expect that

$$G \in \mathcal{D}_{\mathsf{P}}(\mathcal{G}_{\mathcal{A}}) \quad \textit{iff} \quad G \in \mathcal{R}(\mathcal{G}_{\mathcal{A}}) \wedge \neg(\exists q \in P_{rpr}, \exists G' \in \mathcal{R}(\mathcal{G}_{\mathcal{A}}) : G \Rightarrow_q G')$$

As for the case of programmable dynamism, repairing dynamism allows for the formulation of the following two questions:

- the specification is complete. This reduces to prove that $G \in \mathcal{D}_{\mathsf{P}}(\mathcal{G}_{\mathcal{A}})$ implies $G \in \mathcal{R}(\mathcal{G}_{\mathcal{A}}) \wedge \neg(\exists q \in P_{rpr}, \exists G' \in \mathcal{R}(\mathcal{G}_{\mathcal{A}}) : G \Rightarrow_q G')$.
- the specification is correct. This corresponds to prove $G \in \mathcal{R}(\mathcal{G}_{\mathcal{A}}) \wedge \neg(\exists q \in P_{rpr}, \exists G' \in \mathcal{R}(\mathcal{G}_{\mathcal{A}}) : G \Rightarrow_q G')$ implies $G \in \mathcal{D}_{\mathsf{P}}(\mathcal{G}_{\mathcal{A}})$.

In addition, this kind of dynamism naturally poses the question of whether reparing rules are adequate, i.e., whether the set of reparing rules assures that for any configuration that is reachable but not desirable there exists a sequence of repairing rules that moves the configuration to a desirable one. Formally,

- If $G \in \mathcal{R}(\mathcal{G}_{\mathcal{A}}) \wedge (\exists q \in P_{rpr}, \exists G' \in \mathcal{R}(\mathcal{G}_{\mathcal{A}}) : G \Rightarrow_q G')$ then $G \Rightarrow_{q_0} G_1 \Rightarrow_{q_1} \ldots \Rightarrow_{q_n} G_n$ with $G_n \in \mathcal{D}_{\mathsf{P}}(\mathcal{G}_{\mathcal{A}})$ and $\{q_0, \ldots, q_n\} \in P_{rpr}$.

3.3 Ad-hoc dynamism

Roughly ad-hoc dynamism allows the architecture to evolve freely by adding and removing components and connectors without any restriction. The typed grammar corresponding to ad-hoc DSA should therefore exploit a fully general type graph that contains an infinite number of hyperarcs $\mathsf{component}_i$ and $\mathsf{connector}_j$ (see Figure 5), one for every natural $i, j \in \mathbb{N}$. Any hyperarc $\mathsf{component}_i$ ($\mathsf{connector}_j$) stands for the type of all connectors that expose exactly i ports (respectively, j roles). For simplicity, we define all ports as having the same type (otherwise the type graph should

be extended, by adding an infinite number of nodes, to represent every possible port type). Similarly, the set of production is infinite as it must allow for adding/ removing any kind of components and connectors. This leaves little space for verification issues, as the only guarantees given by ad-hoc dynamicity is that reached graphs are software configurations.

3.4 Constructible dynamism

Constructible DSAs are similar to ad-hoc DSAs but here rewriting productions are not the free combination of basic primitives: they are full-fledged programs written in some specific language. The main difference w.r.t. ad-hoc DSA is that a constructible dynamic architecture is mostly characterised by the specific programming language allowed for defining the reconfiguration programs that can manage the evolution. Generally speaking, constructible dynamism provides a very weak notion of desirable configurations, and hence verification aspects are almost meaningless when assuming autonomous reconfiguration (likewise ad-hoc dynamism). However, the situation is slightly different when considering reconfigurations controlled externally (see discussion in Section 5).

4 Automotive Software System

In order to illustrate the main forms of dynamism, we introduce the following scenario borrowed from the European Project SENSORIA.

4.1 Automotive Case Study: Overview

Much of the cost of research and development in vehicle production are associated to automotive software. Today vehicles are equipped with a multitude of sensors and actuators that provide different services, like ABS and vehicle stabilization systems, that assist people to drive safer. Thanks to current mobile technology, vehicles have the possibility to connect to the telephone and internet infrastructures. This has given birth to a variety of new services into the automotive domain. Communication in AS systems may involve *communication* that takes place inside a vehicle (*intra-vehicle)*, connection to vehicles in the vicinity (*inter-vehicle)*, or interaction with the environment, for example through an Internet gateway (*vehicle-environment)*. Our scenario will focus only the last two kinds of communications.

4.2 Car Assistance Scenario

Consider a vehicle subscribed to an assistance service. Due to a collision, the airbag of the car is inflated, which causes the automatic generation of a message destined to the *accident assistant server*. The message can be transmitted through near vehicles until reaching the server (preferred method) or directly to the server. The message will be eventually delivered to the assistance server, which will coordinate the assistance.

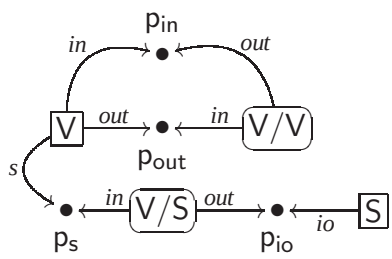

Figure 6. Architectural Style of the AS System

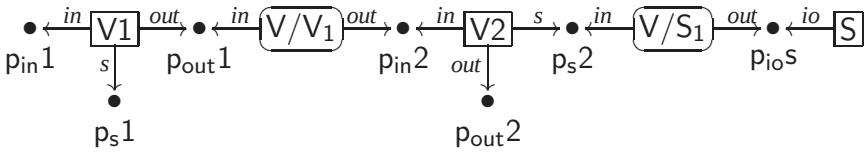

Figure 7. An instance of the AS style.

We will model the scenario by using two different kinds of components:

- **Vehicle (V)**: a component responsible for transmitting messages destined to the assistant server. A vehicle component has three associated ports: p_{in} for receiving a message from another vehicle, p_{out} for sending messages to the next vehicle, and p_s for communicating directly with the server.
- **Accident Assistant Server (S)**: a component that handles help requests. Its unique port p_{io} is used for sending/receiving information to/from vehicles.

Components are connected by using the following connector types:

- **Vehicle to Vehicle communication (V/V)**: a connector used for mediating the communication between two vehicles.
- **Vehicle to Accident Assistant Server communication (V/S)**: a connector used for supporting the interaction between a vehicle and a server.

Figure 6 shows the architectural style of the AS system, while Figure 7 depicts an instance consisting of two vehicles (V1 and V2) and one server (S).

4.3 Programmed Dynamism

We will use a programmable architecture for specifying the way in which the AS system keeps the communication structure among the components. We define the corresponding graph grammar $\mathcal{G}_{\mathcal{A}} = \langle T, G_{in}, P \rangle$, where the architectural style T is depicted in Figure 6, and the initial configuration G_{in} consists in the T-typed graph containing a unique component of type server S and its associated port (i.e., a node of type $\mathsf{port}_{\mathsf{io}}$). The set P contains the productions that model the arrival/departure of a vehicle into/from the area covered by a server, a car that is getting close/far to/from another car, and a car that takes over another one. For space limitations, we describe just three productions. Figure 8 stands for the case of vehicle entering into the area covered by a server. This is modeled by creating a new component of type V and its associated ports, and a new connector between the car and the server.

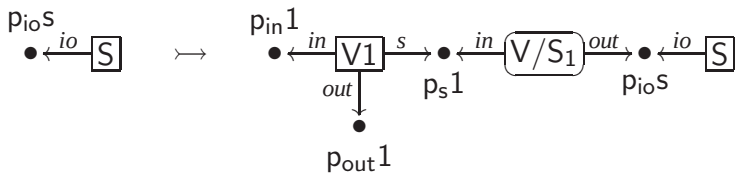

Figure 8. New Vehicle connected to the Server.

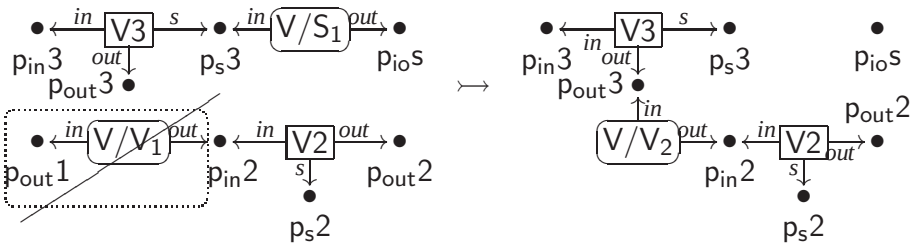

Figure 9. Vehicles approximation.

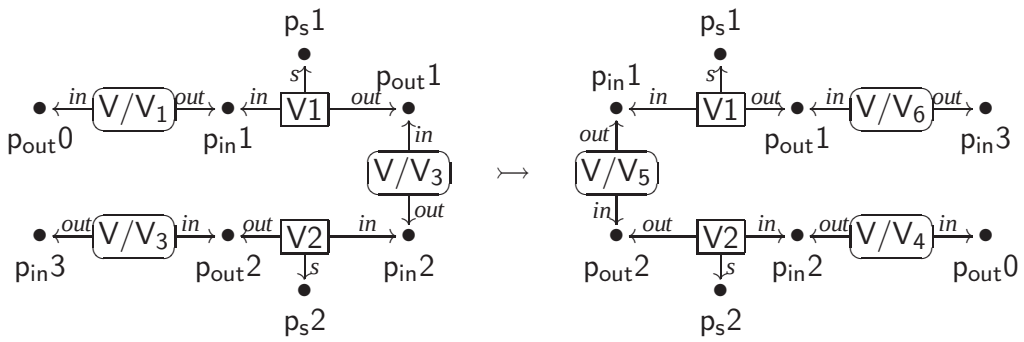

Figure 10. One vehicle takes over another one.

Figure 9 describes the case of a car connected to the server (V3) that gets closer to another vehicle (V2) that is at the end of a queue (note V2 has no connectors attached to its input port). In this case the connection of V3 to the server is removed and a new connector between V3 and V2 is created. Finally, Figure 10 depicts a production that handles the case in which a vehicle takes over another one. (Note the rule assumes the vehicles V1 and V2 to be in the middle of a queue, since they are connected to other vehicles. The specification includes three additional rules to handle the cases in which the take over involves cars at the ends of the queue.)

The set of desirable configurations for the AS system consists of all the configurations in which each vehicle has a unique, acyclic communication path with the unique server, and each vehicle port has attached at most one connector.

4.4 Repairing dynamism

The following example shows the use of a repairing architecture for modelling the fact that the communication between vehicles is not reliable and can be lost, but in such cases the architecture should repair itself in order to provide unconnected components with a link to a server. We defined a graph grammar $\mathcal{G}_{\mathcal{A}} = \langle T, G_{in}, P \rangle$ in which the set of productions is divided into three different sets,

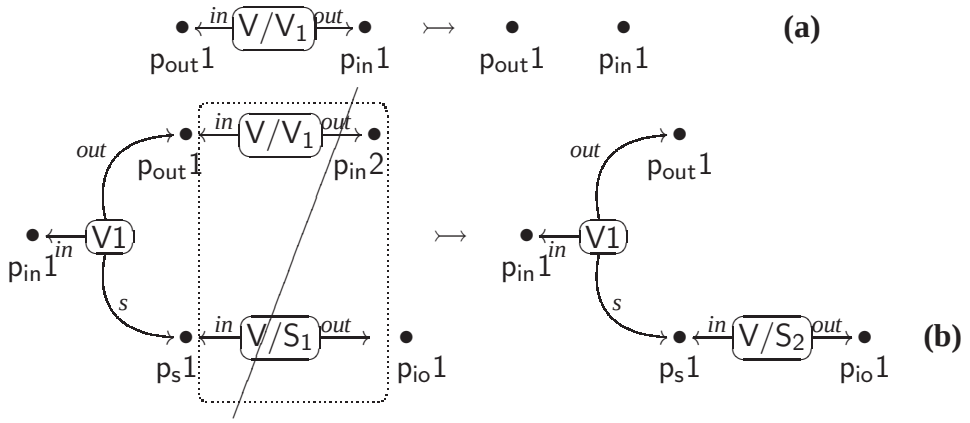

Figure 11. (a)Lost of connectivity. (b) Repairing.

i.e., $P = P_{pgm} \cup P_{env} \cup P_{rpr}$. In our scenario P_{pgm} contains the same productions as defined in Programmed Dynamism (see 4.3), P_{env} contains the unique production shown in Figure 11(a), which models the loss of connectivity between vehicles (i.e., the removal of a connector V/V_1), and repairing production is in Figure 11(b). Such rule states that whenever a vehicle without outcoming connections is found (note the left-hand-side of the rule requires the absence of connections on ports $\mathsf{p}_{\mathsf{out}}1$ and $\mathsf{p}_{\mathsf{s}}1$), then the vehicle should be connected directly to a server.

As for the case of programmable DSAs, in desirable configurations all vehicles have a connection to the server. Nevertheless, the repairing grammar takes into account the cases in which the configuration may not satisfy this condition (some vehicles may not be connected to a server). However, these cases match with the left-hand-side of the repairing rule, and hence, they can be repaired by the system.

5 Constrained and Self dynamism

Other aspects that are, to some extent, orthogonal to the approaches characterised in Section 3 are: (i) whether the application of a transformation rule can take place at any moment or not, and (ii) whether changes are fired internally by the system or activated externally. The first aspect is usually refered to as *constrained vs unconstrained dynamism*, while the second is qualified as *self* vs *external* reconfiguration.

5.1 Unconstrained vs Constrained dynamism

Basically, constrained dynamism refers to the fact that a change may occur only after pre-defined constraints are satisfied. Such constraints may be (i) the configuration topology, e.g., when components are not connected in a specific, or (ii) the state of a component, e.g., when a component enters into the quiescent state. Topological constraints are naturally modelled by both positive and negative application conditions of graph productions. Hence, topological constrained dynamism may be characterised by a graph grammar whose productions have some contexts (either positive or negative). Differently, constraints related to particular states of

components have not an immediate counterpart in our proposal (since our framework does not describe component states). Nevertheless, they can be encoded by thinking about different states of components as different types of hyperedges. In this way, the change of a component state s into s' is represented as the rewrite that removes the hyperarc denoting the component in state s and adds a new hyperarc of type s' with attachments analogous to those of the removed arc. In this case, the fact that the grammar describes a dynamism constrained on the state of some components is hidden by the encoding. Another possibility is to use of attributed graph grammars [18] for equipping components with attributes describing their states.

Unconstrained dynamism refers to the fact that transformations can be applied at any moment. The graph grammar counterpart is the fact that productions have no associated constraints or application conditions, being, in some sense, context free, because they either produce or consume arcs but they do not read them.

5.2 *Self dynamism*

Usually, some kind of dynamisms (like programmed and repairing) are also qualified as "self", meaning that the changes are initiated by the system itself and not by an external agent. We map the notion of self and external dynamism to particular features of the rewrite system. As a starting point we discuss some alternative ways for chosing a particular reconfiguration in a DSA, as proposed in [7].

- *External*: The reconfiguration rule is selected by an external source. This option resembles the external choice of process calculi, in which the branch of computation to be selected is indicated by the context of process. In this sense, we can interpret a reduction of the form $G \Rightarrow_p G'$ as the fact that the environment selects the application of the production p.
- *Autonomous*: The system selects one of all the applicable transformations in a non-deterministic way. This corresponds to the notion of internal choices in process calculi. Accordingly, we may represent such reductions by hiding the actual name of the applied rule. That is, a rewriting step $G \Rightarrow_p G'$ in which p is autonomous can be represented as $G \Rightarrow_\tau G'$, where τ stands for a hidden change.
- *Pre-defined*: Pre-defined selection is a special case of autonomous choice, in which the system selects in a pre-defined way the appropriate transformation to apply from the set of available ones. In this case, the choice is completely deterministic (like a conditional choice if - then - else - of process calculi). This can be mapped into graph grammars as the definition of priorities in the selection of productions to be applied. As shown in [11], application conditions can be used as priorities for restricting the order in which rules are applied.

Let $\mathcal{G} = \langle T, G_{in}, P_{ext} \cup P_{self} \rangle$ be a grammar, where P_{ext} stands for the set of all reconfigurations that are controlled by the environment, while P_{self} contains all the autonomous productions. We say $\mathcal{G}_{\mathcal{A}}$ has (i) self dynamism if $P_{ext} = \emptyset$, (ii) external dynamism if $P_{self} = \emptyset$, or (iii) mixed dynamism otherwise. Assuming that all rewriting steps $G \Rightarrow_p G'$ are written $G \Rightarrow_\tau G'$ when $p \in P_{self}$, we define the

Dynamicity	References	Correctness	Completeness	Adequacy
Programmed	[4,8,13,16,24]	+	+	-
Repairing	[1,9,10,21,22]	+	+	+
Ad hoc	[6,8,25]	-	-	-
Constructible	[2,20]	-/+	-/+	-

Figure 12. Classification summary

following sets associated to the grammar $\mathcal{G} = \langle T, G_{in}, P_{ext} \cup P_{self} \rangle$:

- The set $\mathcal{S}(\mathcal{G})$ of autonomous or self reconfigurations, i.e., the set of all configurations reachable by applying autonomous changes is: $\mathcal{S}(\mathcal{G}) = \{G \mid G_{in} \Rightarrow_{\tau^*} G\}$.
- The set $\mathcal{E}_c(\mathcal{G})$ of reconfigurations associated to an external sequence $c = p_1 \ldots p_n$ of commands: $\mathcal{E}_c(\mathcal{G}) = \{G \mid G_{in} \Rightarrow_{c'} G \wedge c' = \tau^*, p_1, \tau^*, \ldots, \tau^*, p_n, \tau^*\}$. Note $\mathcal{E}_c(\mathcal{G})$ contains all the configurations reachable from the initial configuration by applying the sequence c of external chosen rules interleaved with the application of zero or more autonomous reconfigurations.

Clearly, $\mathcal{S}(\mathcal{G})$ and $\mathcal{E}_c(\mathcal{G})$ are subsets of $\mathcal{R}(\mathcal{G})$. Hence, we can proceed as in Section 3, and formulate some verification problems. In particular, we can specialise the problem $\mathcal{R}(\mathcal{G}) \subseteq \mathcal{D}_{\mathsf{P}}(\mathcal{G})$ to either $\mathcal{S}(\mathcal{G}) \subseteq \mathcal{D}_{\mathsf{P}}(\mathcal{G})$ or $\mathcal{E}_c(\mathcal{G}) \subseteq \mathcal{D}_{\mathsf{P}}(\mathcal{G})$. The last relation is particular interesting when considering ad-hoc or constructible dynamism. In this case, it is possible to check whether a particular reconfiguration program may produce acceptable configurations.

6 Final Remarks

In this work we have characterised different aspects of dynamic reconfiguration as particular features of graph rewriting systems. By taking advantage of this framework, we have distilled whether such kinds of dynamisms allow for posing typical questions about the completeness and correctness of the architectural specification. Figure 12 summarises the conclusions for the different types of dynamisms.

As mentioned in Section 3, given a characterization of all desirable configurations of a programmable architecture, e.g., by defining a property P that should hold in every configuration, then it would be possible to prove whether the architectural specification is correct (by showing that P holds in every reachable configuration) and complete (by proving any configuration satisfying P is reachable). Correctness and completeness properties could be associated to repairing dynamism. But, differently from programmed dynamism, some reachable configurations of a repairing architecture may be non desirable. Instead, those configurations should be transformed into a desirable one by using repairing rules. The main idea is that undesirable configurations are characterized as those reachable configurations in which some repairing rule is applicable. Then, correctness and completeness prop-

erties involve those reachable configurations that are desirable. Such questions are meaningless for ad hoc dynamicity, where every configuration is potentially reachable. Analogously for constructible dynamism, even if some kind of weak analysis could be performed in this case. For instance, to prove that particular configurations are not reachable when the reconfiguration language forbid some kind of programs.

Actually, the above characterization corresponds to the case in which transformations are all autonomous, i.e., when we assume self dynamism. When external dynamism is considered, also correctness and completeness properties over ad hoc and constructible architectures can be formulated. For instance, given a particular (set of) desirable configuration(s) it can be proved whether a particular transformation or configuration program selected by a programmer produces a desirable configuration. Even more interesting is the case in which mixed dynamism is considered. Assume an ad hoc architecture where some productions are considered external and others autonomous or self. In this case, external transformations account for the reconfigurations activated by a user, while autonomous transformations model the actual program that performs the transformation (a kind of scripting). In this case, it would be possible to check whether a particular script produces a correct configuration when it is applied over a specific configuration.

In this paper we have identified classes of properties that can be naturally associated some kinds of dynamicities. Next work will approach the problem of verifying such properties over graph grammar specifications. In particular, we have in mind to use Alloy [14,15] for attempting this task and we are going to concentrate our efforts on proving properties associated to each kind of DSA.

References

[1] R. Allen, R. Douence, and D. Garlan. Specifying and analyzing dynamic software architectures. *Proc. of FASE'98*, vol. 1382 of *LNCS*, pp. 21–37. Springer, 1998.

[2] J. Andersson. Issues in dynamic software architectures. *Proc. of ISAW4*, pp. 111–114, June 2000.

[3] L. Baresi, R. Heckel, S. Thöne, and D. Varró. Modeling and validation of service-oriented architectures: application vs. style. *Proc. of ESEC/FSE'03*, pp. 68–77. ACM, 2003.

[4] L. Baresi, R. Heckel, S. Thöne, and D. Varró. Style-based refinement of dynamic software architectures. *Proc. of WICSA'04*, pp. 155–166. IEEE, 2004.

[5] G. Berry and G. Boudol. The chemical abstract machine. *Theoret. Comput. Sci.*, 96(1):217–248, 1992.

[6] C. Bidan, V. Issarny, T. Saridakis, and A. Zarras. A dynamic reconfiguration service for corba. *Proc. of ICCDS'98*, pp. 35–42, 1998.

[7] J. Bradbury, J. Cordy, J. Dingel, and M. Wermelinger. A survey of self-management in dynamic software architecture specifications. *Proc. of WOSS'04*, pp. 28–33. ACM, 2004.

[8] M. Endler. A language for implementing generic dynamic reconfigurations of distributed programs. *Proc. of BSCN'94*, pp. 175–187, 1994.

[9] D. Garlan and B. Schmerl. Model-based adaptation for self-healing systems. *Proc. of WOSS'02*, pp. 27–32. ACM, 2002.

[10] I. Georgiadis, J. Magee, and J. Kramer. Self-organising software architectures for distributed systems. *Proc. of WOSS'02*, pp. 33–38. ACM, 2002.

[11] A. Habel, R. Heckel, and G. Taentzer. Graph grammars with negative application conditions. *Fundam. Inform.*, 26(3/4):287–313, 1996.

[12] D. Hirsch, P. Inverardi, and U. Montanari. Graph grammars and constraint solving for software architecture styles. *Proc. of ISAW'98*, pp. 69–72. ACM, 1998.

[13] D. Hirsch, P. Inverardi, and U. Montanari. Reconfiguration of software architecture styles with name mobility. *Proc. of COORDINATION'00*, vol. 1906 of *LNCS*, pp. 148–163. Springer, 2000.

[14] D. Jackson. Alloy: a lightweight object modelling notation. *ACM Trans. Softw. Eng. Methodol.*, 11(2):256–290, 2002.

[15] D. Jackson. *Software Abstractions: Logic, Language, and Analysis*. The MIT Press, 2006.

[16] D. Le Métayer. Describing software architecture styles using graph grammars. *IEEE Trans. Software Eng.*, 24(7):521–533, 1998.

[17] M. Löwe. Algebraic approach to single-pushout graph transformation. *Theoret. Comput. Sci.*, 109(1-2):181–224, 1993.

[18] M. Löwe, M. Korff, and A. Waner. An algebraic framework for the transformation of attributed graphs. *Term Graph Rewriting: Theory and Practice*, pp. 185–199. John Wiley & Sons Ltd, 1993.

[19] N. Medvidovic and R. Taylor. A classification and comparison framework for software architecture description languages. *IEEE Trans. Software Eng.*, 26(1):70–93, 2000.

[20] P. Oreizy. Issues in the runtime modification of software architecture. Technical Report UCI-ICS-96-35, University of California, 1996.

[21] P. Oreizy, M. Gorlick, R. Taylor, D. Heimhigner, G. Johnson, N. Medvidovic, A. Quilici, D. Rosenblum, and A. Wolf. An architecture-based approach to self-adaptive software. *IEEE Intelligent Systems*, 14(3):54–62, 1999.

[22] B. Schmerl and D. Garlan. Exploiting architectural design knowledge to support self-repairing systems. *Proc. of SEKE'02*, pp. 241–248. ACM, 2002.

[23] M. Shaw and D. Garlan. *Software Architecture: Perspectives on An emerging Discipline*. Prentice Hall, 1996.

[24] M. Wermelinger. Towards a chemical model for software architecture reconfiguration. *IEE Proc. - Software*, 145(5):130–136, 1998.

[25] A. Young and J. Magee. A flexible approach to evolution of reconfigurable systems. *Proc. of IWCDS'92*, IEE, pp. 152–163, 1992.