Semantics of Transient Connectors in Rewriting Logic

José Luiz Fiadeiro
Departamento de Informática, Faculdade de Ciências
Universidade de Lisboa, Campo Grande, 1700 Lisboa, Portugal
E-mail: llf@di.fc.ul.pt

Michel Wermelinger
Departamento de Informática, Faculdade de Ciências e Tecnologia
Universidade Nova de Lisboa, 2825 Monte da Caparica, Portugal
E-mail: mw@di.fct.unl.pt

José Meseguer
Computer Science Laboratory, SRI International
333 Ravenswood Ave, Menlo Park, CA 94025, USA
E-mail: meseguer@csl.sri.com

Abstract

Research in Software Architectures has put forward the concept of connector to express complex relationships between system components, thus facilitating the separation of coordination from computation. A system can then be understood, at a given level of abstraction, in terms of its components and the connectors that establish how they interact. However, for systems in which many interconnections exist between their components, the architectures themselves may become very complex due to the high number of connectors in place. This is especially true in the context of mobile systems in which the interconnections are, by nature, transient in the sense that, at a given instant of time, only a subset of the potential connectors are actually effective. In this paper, we formalise a notion of transient connector that allows, at any given moment, for the architecture to depict only the connectors that are active and, in this way, capture the dynamics of architectures themselves. Our approach is based on the use of COMMUNITY, a UNITY-like program design language that has a semantics in Category Theory, and rewriting logic as a means of capturing the dynamic aspects of connectors.

1 Introduction

In a previous paper [16] we have argued in favour of a disciplined approach to mobility through the use of connectors (in the sense of software architectures). The idea is that mobility within a system can be characterised by the transient nature of the interconnections that exist between the components of the system. Because, from an architectural point of view, such interconnections are best captured through the use of connectors [15], changes in the interconnections should be also captured at the level of

the connectors that are in place.

For that purpose, we defined connectors which, through guarded actions, were able to set or reset interconnections between components according to given conditions of applicability (coded in the guards of the actions). However, because the dynamic behaviour of the system may require a considerable number of different situations in which such interconnections should apply, the architecture of the system may get cluttered by a high number of connectors, even if, at each given time, only a few of the applicability conditions hold.

To circumvent this problem, we suggested in [17] the concept of a transient connector in the sense of a connector with an associated condition on the state of its roles that determines the situations in which it applies. The idea is that a connector does not need to be permanently part of an architecture, but is added and removed according to its applicability condition. This can be seen as a restricted form of dynamic architectures in which the evolution of the architecture is determined by well defined operations of addition and removal of connectors that are to be performed in well determined states of the underlying system.

Our purpose in this paper is to expand the original motivation and further develop the notion of transient connector in the context of dynamic architectures, namely by providing a well-defined mathematical semantics through which the evolution of the architecture can be inferred and reasoned about. Capitalising on previous work on the formalisation of architectural connectors in general [5], and connectors for mobile systems in particular [16], we use Category Theory to represent software architectures. For modelling the dynamic aspects of architectures, we use Rewriting Logic [10], a formalism that has already been applied to the formalisation of several architectural aspects of systems, e.g. [11, 12]. We illustrate the approach with a connector for partial synchronisation of actions written in COMMUNITY [6].

2 Preliminaries

Our example is inspired in the luggage distribution system also used to illustrate Mobile UNITY [14]. One or more carts move on a N units long track with the shape

A cart advances one unit at each step in the direction shown by the arrows. The i-th cart starts from a unit determined by an injective function $start$ of i. Carts are continuously moving around the circuit. Their movement must be synchronised in such a way that no collisions occur at the crossing.

COMMUNITY [6] is a program design language based on UNITY [2] and IP [7]. In this paper we only consider a subset of the full language. For our purposes, a COMMUNITY program consists of a set of typed attributes, a boolean expression to be satisfied by the initial values of the attributes, and a set of actions, each of the form *name*: [*guard* $\rightarrow$ *assignment(s)*]. The empty set of assignments is denoted by `skip`. Action names act as *rendez-vous* points for program synchronisation. At each step, one of the actions is selected and, if its guard is true, its assignments are executed simultaneously. To be more precise, syntactically a program has the form

```
program P is
var    V
read   R
init   I
do     a1: [g1 →  v11 := exp11 ‖ v21 := exp21 ‖ ...]
[]     a2: [g2 → v12 := exp12 ‖ ...]
[]     ...
```

where R are external attributes (i.e., the program may not change their values), V are the local attributes, I is the initialisation condition on V, a_i are the actions with boolean expressions g_i over $V \cup R$, $v_{ij} \in V$, and exp_{ij} expressions over $V \cup R$.

The following program describes the behaviour of the i-th cart.

```
program Cart_i is
var    l : int
init   l = start(i)
do     move: [true → l := (l + 1) mod N]
```

We henceforth omit the "mod N" operation and the action guards whenever they are "true".

A morphism from a program P to a program P' states that P is a component of the system P' and, as shown in [6], captures the notion of program superposition [2, 7]. Mathematically speaking, the morphism maps each attribute of P into a attribute of P' of the same type, and it maps each action name g of P into a (possible empty) set of action names $\{g'_1, \ldots, g'_n\}$ of P' [16]. Those actions correspond to the different possible behaviours of g within the system P'. These different behaviours usually result from synchronisations between g and other actions of other components of P'. Thus each action g'_i must preserve the functionality of g, possibly adding more things specific to other components of P'. In particular, the guard of g'_i must not be weaker than the guard of g, and the assignments of g must be contained in g'_i.

It can be proved that COMMUNITY programs and their morphisms form a category in which every finite diagram has a colimit, which, by definition, is the minimal program that contains all programs in the diagram. Thus the diagram specifies the architecture and the colimit represents the resulting system. Since the proof of the existence of a colimit is constructive, the architecture can be "compiled" into a single program that simulates the execution of the overall system.

3 Transient Connectors

A n-ary connector consists of n roles R_i and one glue G stating the interaction between the roles. These act as "formal parameters", restricting which components may be linked together through the connector. Thus, the roles may contain attributes and actions which are not used for the interaction specification.

Applying these ideas to connectors [5], for each role R_i there must be a channel C_i together with morphisms $\gamma_i : C_i \to G$ and $\rho_i : C_i \to R_i$ stating which attributes and actions of R_i are used in the interaction specification, i.e., the glue. As channels just establish the required relationships between action names, their actions are always of the form a: [true → `skip`], and thus in this paper we use the abbreviated notation

```
channel C is
read   v1 : T1; ...
do     a1, a2, ...
```

The categorical framework also allows one to make precise when an n-ary connector can be applied to components $P_1,\ldots,P_n$, namely when morphisms $\iota_i : R_i \to P_i$ exist. This corresponds to the intuition that the "actual arguments" (i.e., the components) must instantiate the "formal parameters" (i.e., the roles). As an illustration, an instantiated binary connector has the diagram

$$P_1 \xleftarrow{\iota_1} R_1 \xleftarrow{\rho_1} C_1 \xrightarrow{\gamma_1} G \xleftarrow{\gamma_2} C_2 \xrightarrow{\rho_2} R_2 \xrightarrow{\iota_2} P_2$$

We proposed in [17] the use of transient connectors as consisting of a pair $\langle I, T\rangle$ where T is a connector and I is a boolean expression, called interaction condition, over the attributes of its roles.

Returning to our example, assume that two carts are approaching the crossing and one of them is nearer to it. To avoid a collision it is sufficient to force the nearest cart to move whenever the most distant one does. That can be achieved using an action subsumption connector. Action a subsumes action b if b executes whenever a does. This can be seen as a partial synchronisation mechanism: a is synchronised with b, but b can still execute freely. The connector that establishes this form of interaction is

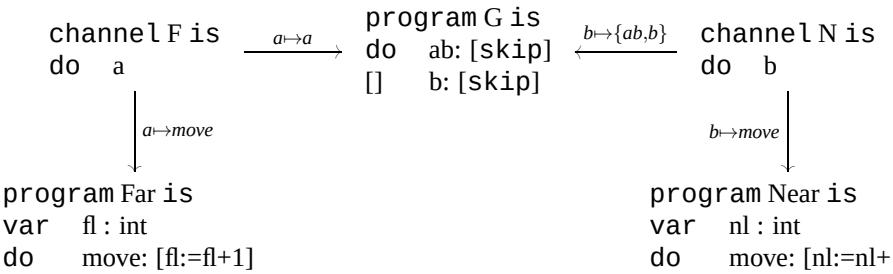

Notice that although the two roles are isomorphic, the binary connector is not symmetric because the glue treats the two actions differently. This is clearly indicated in the glue: "b" may be executed alone at any time, while "a" must co-occur with "b" if the interaction is taking place. Hence, action "a" is the one that we want to connect to the "move" action of the cart that is further away from the crossing, while action "b" is associated to the movement of the nearest cart (the one that will instantiate role "Near").

To complete the example, it remains to show what the interaction condition is, and what system is obtained through role instantiation. Assuming that track units 7 and 28 cross and that movement coordination should start when both carts are at most 3 units away from the crossing, one has

$$I = 0 \leq 7 - nl < 28 - fl \leq 3 \vee 0 \leq 28 - nl < 7 - fl \leq 3$$

The first disjunct treats the case when the nearest cart is moving towards track unit 7 and the other cart is approaching unit 28. The other disjunct handles the opposite case.

However, whereas with permanent connectors the configuration of a system could be represented in a mathematically simple way as a (complex) diagram whose colimit returns the behaviour of the system, with transient connectors we avoid the explosion of connectors that clutters the configuration diagram, but we need to provide a mathematical semantics for them.

4 Transient connectors as rewrite rules

Rewriting logic [8, 9] expresses an essential equivalence between logic and computation in a particularly simple way. Namely, system *states* are in bijective correspondence with *formulas* (modulo whatever structural axioms are satisfied by such formulas: for example, modulo the associativity and commutativity of a certain connective) and concurrent *computations* in a system are in bijective correspondence with *proofs* (modulo appropriate notions of equivalence between computations and between proofs). Given this equivalence between computation and logic, a rewriting logic axiom of the form

$$t \longrightarrow t' \textbf{ if } C$$

has two readings. Computationally, it means that a fragment of a system's state that is an instance of the pattern t can *change* to the corresponding instance of t' concurrently with any other state changes when condition C holds; that is, the computational reading is that of a *local concurrent transition*. Logically, it just means that we can derive the formula t' from the formula t when C holds; that is, the logical reading is that of an *inference rule*.

Rewriting logic is entirely neutral about the structure and properties of the formulas/states t. They are entirely *user-definable* as an algebraic data type satisfying certain equational axioms, so that rewriting deduction takes place *modulo* such axioms. More precisely, a *signature* in rewriting logic is an equational theory (Σ, E), where Σ is an equational signature and E is a set of Σ-equations. Rewriting will operate on equivalence classes of terms modulo E. In this way, rewriting is made free from the syntactic constraints of a term representation and gain a much greater flexibility in deciding what counts as a *data structure*; for example, string rewriting is obtained by imposing an associativity axiom, and multiset rewriting by imposing associativity and commutativity. Of course, standard term rewriting is obtained as the particular case in which the set of equations E is empty. Techniques for rewriting modulo equations have been studied extensively [4] and can be used to implement rewriting modulo many equational theories of interest.

Given a signature (Σ, E), *sentences* of rewriting logic are sequents of the form

$$r : [t]_E \longrightarrow [t']_E \textbf{ if } C,$$

where r is a label, t and t' are Σ-terms possibly involving some variables, $[t]_E$ denotes the equivalence class of the term t modulo the equations E, and C is a condition expressed as a conjunction of equations or sequents of the form $[u_i] \longrightarrow [v_i]$. A *rewrite theory* R is a 4-tuple $R = (\Sigma, E, L, R)$ where Σ is a ranked alphabet of function symbols, E is a set of Σ-equations, L is a set of *labels*, and R is a set of sentences as described above, called *rewrite rules*.

Because of its neutrality with regard to the structure and properties of states and formulas, rewriting logic has good properties as a *semantic framework* [10], in which many different system styles and models of concurrent computation and many different languages can be naturally expressed without any distorting encodings.

For instance, a COMMUNITY program can be represented as a rewrite theory whose signature defines state configurations as sets of pairs $\langle a : T \mid val : v \rangle$ with a a program attribute, T a type, and v a value of type T, and every action a as a rewrite rule:

$$a: \quad \begin{array}{l} \langle a_1 : T_1 \mid val : x_1 \rangle \ldots \langle a_n : T_n \mid val : x_n \rangle \\ \longrightarrow \langle a_1 : T_1 \mid val : exp_1(x_1, \ldots, x_n) \rangle \ldots \langle a_n : T_n \mid val : exp_n(x_1, \ldots, x_n) \rangle \end{array} \textbf{ if } g$$

where g, the guard of the action, is a condition on the x_i. In this case, the equational axioms E modulo which we rewrite are the associativity and commutativity of set union, which is expressed in such rule by empty syntax (juxtaposition).

Graph rewriting has also been represented in rewriting logic [10]. Labelled graphs are axiomatised equationally as an algebraic data type in such a way that graph rewriting becomes rewriting modulo the equations axiomatising the type. Axiomatisations in this spirit include those of Bauderon and Courcelle [1], Corradini and Montanari [3], and Raoult and Voisin [13]. We adopt the axiomatisation given in [10], in which a (labelled) graph is viewed as a set of nodes, and thus graph rewriting is viewed modulo the associativity and commutativity of set union, expressed again with empty syntax.

The labels that interest us for the semantics of transient connectors are pairs (P,s) with P a COMMUNITY program and s a state configuration for P. Edges between nodes labelled (P',s') are labelled with morphisms $f : P \to P'$ such that, for every $\langle a : T \mid val : v\rangle$ in s, $\langle f(a) : T \mid val : f(v)\rangle$ is in s' (modulo the equational axioms E), i.e., morphisms have to respect the state configurations. We shall call such graphs *anchored configurations*.

The idea is to represent an n-ary transient connector defined by $\gamma_i : C_i \to G$ and $\rho_i : C_i \to R_i$ as a conditional graph rewrite rule of the form

$$\begin{array}{ccccccc} & & C_1 & \xrightarrow{\gamma_1} & (G,s) & \xleftarrow{\gamma_n} & C_n \\ (P_1,s_1)\cdots(P_n,s_n) & \longrightarrow & \downarrow{\scriptstyle \rho_1;X\iota_1} & & \cdots & & \downarrow{\scriptstyle \rho_n;X\iota_n} \\ & & (XP_1,s_1) & & & & (XP_n,s_n) \end{array}$$

if $I \wedge X\iota_1 \in morph(R_1,XP_1) \wedge \cdots \wedge X\iota_n \in morph(R_n,XP_n)$

where the XP_i are "variables" that can be instantiated with any programs subject to the conditions imposed by the rule, which are, for each instance P_i, that it admits the corresponding instance of s_i as a valid configuration, and that an instance of $X\iota_i$ be found that is a morphism from the connector's role R_i to the instance P_i (thus making P_i a true instance of the role in the categorical sense as discussed in section 3). Notice that each instance P_i will be connected to the glue via the channel C_i and the morphism that results from the composition of the morphisms that connect the channel to the role, as given by the connector, and the instance of $X\iota_i$ that establishes P_i as an instance of R_i. The instances of the state configuration must, of course, satisfy the interaction condition I. Finally, s is the state given by the initialisation condition of the glue G.

For instance, in the case of the cart synchronisation, we would have for the connector defined in section 3 the rewrite rule corresponding to:

$$\begin{array}{l} (XP_1,\langle Xf : int \mid val : f\rangle,XL)(XP_2,\langle Xn : int \mid val : n\rangle,XM) \\ \longrightarrow \\ \qquad\qquad F \xrightarrow{a\mapsto a} (G,nil) \xleftarrow{b\mapsto\{ab,b\}} N \\ \qquad a\mapsto move;X\iota_1 \downarrow \qquad\qquad\qquad\qquad b\mapsto move;X\iota_2 \downarrow \\ (XP_1,\langle Xf : int \mid val : f\rangle,XL) \qquad\qquad (XP_2,\langle Xn : int \mid val : n\rangle,XM) \end{array}$$

if

$0 \le 7-n < 28-f \le 3 \vee 0 \le 28-n < 7-f \le 3 \wedge X\iota_1 \in morph(Far,P_1) \wedge X\iota_2 \in morph(Near,P_2)$

The expression of the rule in rewriting logic, through the use of variables ranging over nodes, programs, morphisms and lists of edges, makes clear that the left hand-side of the rules can be instantiated by any nodes labelled by any programs matching the given state configuration, and with any connectivity to other nodes, subject to the applicability conditions. These include the interaction condition of the programs and the identification of the morphisms that are being used to instantiate the roles. These conditions on the instantiation morphisms are essential to narrow down the scope of the applicability of the rule to programs that actually fit the roles.

Because the left-hand side of the rule is copied to the right hand side, the effect of the application of the rule is to superpose the glue and its connections to the components identified through the left-hand side. The fact that the identifiers of the superposed nodes and edges are new means that the interconnections are, indeed, new and do not interfere with other interconnections that may exist.

Notice that the reverse rewrite rules, removing the application of the connectors, are also necessary when the interaction condition becomes false.

Summarising, the architecture of the system consists of a rewrite theory presentation over the signature that we have outlined above in terms of anchored configurations. The axioms of this rewrite theory presentation are the conditional rewrite rules defined by the connectors. Given an initial anchored configuration of the system, such a rewrite theory presentation provides us with the space of possible evolutions of the system configuration from that state.

5 Concluding Remarks

Transient connectors state explicitly the condition that programs must obey in order to interact according to the way prescribed by the connector. Externalising the interaction condition makes the connectors simpler and allows their reuse under different circumstances. The architectural diagram also becomes simpler (and more intuitive) since it reflects at each point in time just the connections that are in place.

In this paper we have given an operational semantics for transient connectors in rewriting logic. For each connector there are two rules, one to introduce it into the architecture, the other one to remove it. In both cases, the left and right hand sides of the rules are anchored configurations showing the current state of each program, thus allowing the evaluation of the interaction condition of the connector. Programs are written in a UNITY-like language, which also has a semantics in rewriting logic. This allows a uniform representation of both the computational and the architectural levels, showing how they interact, and of their dynamics, showing how they jointly evolve.

There is also an added expressive power in the proposed semantics of architecture that we intend to explore in future work: the ability of actions to be constrained by conditions on the structure of the configuration. Further work that we intend to pursue includes the definition of rewriting strategies for execution in Maude, and the use of logical mechanisms available for reasoning about rewrite theories for reasoning about the evolution of the systems that are subject to transient connectors.

Acknowledgements

This work was partially supported by Fundação para a Ciência e Tecnologia (FCT) through contracts PRAXIS XXI PCEX/P/MAT/46/96 (ACL) and PRAXIS XXI 2/2.1/TIT/1662/95 (SARA). Support from the US Office of Naval Research under contract N00014-96-C-0114, and the National Science Foundation through grant CCR-9633363 are also gratefully acknowledged. This work was produced while José Fiadeiro was on leave at the SRI International with the support of FLAD and FCT (contract BPD 16367/98).

References

[1] M. Bauderon and B. Courcelle. Graph expressions and graph rewriting. *Math. Systems Theory*, 20:83–127, 1987.

[2] K. M. Chandy and J. Misra. *Parallel Program Design: A Foundation*. Addison-Wesley, 1988.

[3] A. Corradini and U. Montanari. An algebra of graphs and graph rewriting. In D. P. et al., editor, *Category Theory and Computer Science*, pages 236–260. Springer LNCS 530, 1991.

[4] N. Dershowitz and J.-P. Jouannaud. Rewrite systems. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Vol. B*, pages 243–320. North-Holland, 1990.

[5] J. L. Fiadeiro and A. Lopes. Semantics of architectural connectors. In *Proceedings of TAPSOFT'97*, volume 1214 of *LNCS*, pages 505–519. Springer-Verlag, 1997.

[6] J. L. Fiadeiro and T. Maibaum. Categorial semantics of parallel program design. *Science of Computer Programming*, 28:111–138, 1997.

[7] N. Francez and I. Forman. *Interacting Processes*. Addison-Wesley, 1996.

[8] J. Meseguer. Rewriting as a unified model of concurrency. In *Proceedings of the Concur'90 Conference, Amsterdam, August 1990*, pages 384–400. Springer LNCS 458, 1990.

[9] J. Meseguer. Conditional rewriting logic as a unified model of concurrency. *Theoretical Computer Science*, 96(1):73–155, 1992.

[10] J. Meseguer. Rewriting logic as a semantic framework for concurrency: a progress report. In *Proceedings of the 7th International Conference on Concurrency Theory*, volume 1119 of *LNCS*, pages 331–372. Springer-Verlag, 1996.

[11] J. Meseguer and C. Talcott. Semantic interoperation of dynamic heterogeneous architectures, 1997. Technical presentations to EDCS Architecture Cluster Meeting, April and July 1997.

[12] J. Meseguer and C. Talcott. Using rewriting logic to interoperate architectural description languages (I and II). Lectures at the Santa Fe and Seattle DARPA-EDCS Workshops, March and July 1997., 1997.

[13] J.-C. Raoult and F. Voisin. Set-theoretic graph rewriting. In H.-J. Schneider and H. Ehrig, editors, *Graph Transformations in Computer Science*, pages 312–325. Springer LNCS 776, 1994.

[14] G.-C. Roman, P. J. McCann, and J. Y. Plun. Mobile UNITY: Reasoning and specification in mobile computing. *ACM TOSEM*, 6(3):250–282, July 1997.

[15] M. Shaw and D. Garlan. *Software Architecture: Perspectives on an Emerging Discipline*. Prentice Hall, 1996.

[16] M. Wermelinger and J. L. Fiadeiro. Connectors for mobile programs. *IEEE Transactions on Software Engineering*, 24(5):331–341, May 1998.

[17] M. Wermelinger and J. L. Fiadeiro. Towards an algebra of architectural connectors: a case study on synchronization for mobility. In *Proceedings of the Ninth International Workshop on Software Specification and Design*, pages 135–142. IEEE Computer Society Press, 1998.